Measuring Welfare Gains from Online Stores:

Theory and Evidence from the Supreme Court's Wayfair Decision*

Zijian He† Yang Su ‡

Job Market Paper

January 31, 2023

Click here for latest version

Abstract

We study how the rise of e-commerce has reshaped consumer welfare and its distributional implications in the presence of retail oligopoly. Based on new data on shopping receipts, we document consumer heterogeneity in online retailing markets: households living in rural areas and with higher incomes are more likely to shop online. To quantify the welfare effects, we leverage an exogenous tax shock by the Supreme Court's Wayfair Decision to learn about online store substitutability and firm pricing responses. We then develop and estimate a structural demand and supply model focussing on the pet food retail market. The model allows us to decompose the consumer online welfare gains into gains from varieties (9%) and convenience (5%) and gains from pro-competitive effects (3%). We further characterize the distributional effects of the rise of e-commerce and find it has reduced consumption inequality between rural and urban areas but increased consumption inequality between the rich and the poor.

*Zijian thanks Costas Arkolakis, Samuel Kortum, and Katja Seim for their invaluable mentorship and support throughout the project. He also thanks Steven Berry, Sheng Cai, Lorenzo Caliendo, Haoge Chang, Ana Fieler Cecilia, Philip Haile, Charles Hodgson, Justin Leung, Jintaek Song, Michael Sullivan, and seminar participants of the Trade and IO lunch/workshops at Yale University and UEA for helpful discussions. Researcher(s)' own analyses calculated (or derived) based in part on data from Market Track, LLC dba Numerator and marketing databases provided through the Numerator Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Numerator data are those of the researcher(s) and do not reflect the views of Numerator. Numerator is not responsible for and had no role in analyzing or preparing the results reported herein.

†Yale University, Department of Economics, zijian.he@yale.edu

‡University of Chicago, Booth School of Business, ysu4@chicagobooth.edu

1. Introduction

The market share of e-commerce in the U.S. retail sector increased from 5% in 2012 to 15% in 2021.[1] The online retail market is concentrated among only a few firms. Recently, there have been intensive policy debates regarding online market regulations such as taxation and antitrust laws. Before addressing these policy debates, it is critical to answer a fundamental question about e-commerce: how does the rise of e-commerce affect overall consumer welfare and consumption inequality? First, e-commerce may benefit the average consumer by providing them with additional shopping options and saving them travel time. These gains are referred to as gains from variety and convenience, respectively (Dolfen et al. (2019)). At the same time, e-commerce may foster competition, making traditional sellers lower their prices. This pro-competitive effect remains unexplored. Second, the rise of e-commerce may have distributional implications across various geographical and income groups that are important for policy evaluation. For example, if rural areas and high-income households were found to rely more on online shopping than other groups, an online sales tax would function as both a spatial redistribution policy and an income-progressive tax.

When studying the effects of the rise of e-commerce on consumer welfare, there are two empirical challenges: (1) limited data describing individual omnichannel[2] shopping behavior to learn about consumer heterogeneity, and (2) no credible estimates of the online-offline substitution elasticity, which is indispensable when using observed data to infer consumer welfare. We make progress on both fronts by combining novel data from consumer shopping receipts sourced from Numerator as well as an exogenous tax shock, the Supreme Court's South Dakota v. Wayfair, Inc., Overstock.com, Inc., and Newegg, Inc. (2018) decision (hereafter, "Wayfair Decision"), to estimate the substitutability of online stores.

The Numerator data have better online coverage and more accurate tax information than existing data from Nielsen and Comscore. More importantly, they reveal retailer information, allowing us to link online stores to their offline counterparts to study online-offline substitution. We can also use the data to measure online consumption heterogeneity across geographical and income groups. We find that (1) households living in rural areas rely more on online stores (the interquartile range of ZIP online expenditure share is 15%); and (2) conditional on the ZIP,

[1] Source: U.S. Department of Commerce Cb22-23.

[2] We use the term "offline" to refer to bricks-and-mortar retailers in this paper. "Omnichannel" means both online and offline shopping channels.

households with higher incomes are 30% more likely to shop online.

To measure online store substitutability, we leverage the Supreme Court's Wayfair Decision, which causes substantial variations in both the level and timing of sales tax changes across markets for out-of-state retailers. Before the Wayfair Decision, online sellers without a physical presence in one state had no obligation to collect sales tax on out-of-state sales.[3] After the Wayfair Decision, however, sellers were legally obligated to collect local sales tax ranging from 0% to 12.5% depending on the market. Additionally, each state enacted the sales-tax law by setting a state-specific deadline by which online sellers were required to adopt the tax change. We use all of the online receipts in the Numerator data to study the effect of the Wayfair Decision on all sectors. However, we focus on the pet-food retail market for online-offline substitution and equilibrium effects for several reasons: (1) pet food is a well-defined category with large online shares; (2) our data do not cover sales from stores specialized in durable goods (e.g., Best Buy, IKEA); and (3) the pet-food market is one of the sectors that were substantially affected by the Wayfair Decision. We avoid studying the grocery market since it is tax-free and thus not subject to the Wayfair Decision.

We first deploy an event-study design using the staggered adoption of the Wayfair Decision to visualize the average effect on sales and firm pricing responses. We find that tax shocks, on average, decrease the sales of affected online retailers by 5% and have no effect on local online pre-tax prices. Our findings suggest that online retailers tend to set uniform prices and do not tailor their prices to local shocks. For the indirectly treated retailers, we find small increases in their local prices and no assortment adjustment in response to the Wayfair Decision. Next, we use both the timing and the level of tax shocks to document reduced-form evidence of online store substitutability. We regress relative sales on relative (after-tax) prices, using the tax shocks as price instruments. We find stores belonging to the online channel and stores belonging to the same retailer are more substitutable than others.

Motivated by our findings on consumer heterogeneity and online store substitutability, we build a demand and supply model of the pet-food retail market to study the equilibrium effect on firm pricing and consumer welfare. On the demand side, we model consumer shopping store choices based on prices, taxes, distance, and quality, allowing for flexible substitution patterns. We use a generalized extreme value (GEV) specification to capture the idea that online stores may be close substitutes for other online stores and that stores belonging to the same

[3] The Supreme Court's decision in Quill Corp. v. North Dakota (1992) established this precedent.

retailer may also be close substitutes. Our specification yields closed-form market shares and price indexes. Our new retailer-level price index generalizes the canonical constant elasticity of substitution (CES) price index by adding qualities and spatial frictions with rich substitution elasticities. On the supply side, we model firms incurring different logistics costs to fulfill online and offline orders and assume they engage in a Bertrand-Nash pricing game.

The model has rich substitution patterns, spatial frictions, consumer heterogeneity, and unobserved demand and supply shocks. We identify the parameters by exploiting variation in tax shocks as well as spatial variation in consumers' choice sets and firms' fulfillment centers. We use the instrumental variables estimation using the generalized method of moments (IV-GMM) procedure from the empirical industrial-organization literature (see Berry et al. (1995)) for our demand and supply estimation. Our demand estimates are consistent with the reduced-form results. We find that lower-income people are more price-sensitive and distance-averse. In addition, stores belonging to the same channel and retailer are more substitutable. The supply estimates show that online logistics costs are 40% higher than offline logistics costs, conditional on the same distance.

The demand and supply model is a laboratory to evaluate many counterfactual online policies. As a first exercise, we calculate the welfare effects of the Wayfair Decision itself, holding prices fixed. The change in the online sales tax regime decreases consumer welfare by an average of 1.6% (in terms of a household's pet-food budget) and hurts the rich more than the poor. For the second exercise, we decompose the welfare change due to the rise of e-commerce into different channels considering firm endogenous price-setting. Starting from the 2021 economy, we first assume online shopping incurs the travel cost of consumers' average distance to offline stores. We call the welfare change in this step gains from convenience. We find that consumers' price index (inverse of welfare) increases by 5%. Then we remove all the online stores but hold all remaining retailer prices fixed to measure gains from online variety. We find welfare decreases by an additional 9%. Finally, we allow all remaining firms to optimally adjust their prices to quantify the pro-competitive effect. We find that the offline sellers increase their prices by an average of 13%, which leads to a further 3% increase in the price index. Overall, the online welfare gains in the pet-food market are 17%. Finally, we characterize the distributional effect of the rise of e-commerce and find it reduces consumption inequality between rural and urban areas but increases consumption inequality between the rich and the poor.

We make the following contributions to the literature. First, we offer a new view of con-

sumer welfare based on the retail landscape, which depends on distance, prices, quality, and the presence of stores. Our work expands on the view of traditional consumer welfare measurement based solely on products. In the consumption goods sector, Handbury and Weinstein (2015), Handbury (2021), and Diamond and Moretti (2021) study consumer welfare based on the product level and do not consider how accessibility to retailer affects consumer welfare. In the service sector, Davis et al. (2019), Miyauchi et al. (2021), Seo and Oh (2022), and Couture (2013) study how distance to restaurants affects consumer welfare, but they do not consider retail oligopolies and their price-setting. The economic-geography literature highlights how market access determines consumer welfare (e.g., Allen and Arkolakis (2014), Donaldson and Hornbeck (2016)). We show that retail chains absorb the shipping costs into their uniform pricing. These papers do not consider how the online market has changed consumer welfare. Of the recent studies on e-commerce that find evidence of consumer benefit through various channels, such as Dolfen et al. (2019), Huang and Bronnenberg (2021), Couture et al. (2021), and Jo et al. (2019), Dolfen et al. (2019) is the closest to our paper. They study the gains from convenience and variety under a CES monopolistic framework. We add gains from the pro-competitive effect in retail oligopoly and study the distributional effect.

We also contribute to the literature by using institutional tax variation to identify demand and supply. Zoutman et al. (2018) formally established this idea in log-linear systems, and Dearing (2022) generalized this idea to general demand and supply systems. In the recent trade literature, Fajgelbaum et al. (2020) uses tariff changes to identify nested-CES demand and supply. We apply this strategy in the urban setting in the presence of retail oligopolies. Some existing literature uses the cross-sectional variation in tax rates before the Wayfair Decision to learn about the online market (e.g., Einav et al. (2014), Houde et al. (2021), and Hossain (2022)). We complement this strategy by incorporating tax shocks, thereby eliminating potential concerns about correlated demand shifters and tax rates. Most importantly, since our data allow us to observe consumer omnichannel behavior and estimate demand and supply, we can characterize the market equilibrium and analyze consumer welfare.

Lastly, our study is related to the spatial retail literature. First, we show that even stores belonging to the same firm are imperfect substitutes. Assuming perfect substitutes among retail chains understates the consumer welfare gains from branching (e.g., Rossi-Hansberg et al. (2020), Huang and Bronnenberg (2021)). Second, we offer a comprehensive view of the importance of geography given the presence of e-commerce. Due to uniform pricing, online demand

is less subject to geographical variations (Fan et al. (2018), Goolsbee (2000)), yet space still matters. Our estimates suggest logistics costs of fulfilling online orders in the form of parcel shipping are higher than those of fulfilling offline orders by a full truckload (FTL). Previous literature such as Seim (2006) and Jia (2008) uses the equilibrium spatial distribution of firms to infer their profit function. We focus on the implications of the spatial distributions of firms on consumer welfare instead. Third, we contribute to the literature on the price responses of big firms (e.g., DellaVigna and Gentzkow (2019), Adams and Williams (2019), Butters et al. (2022)). We have evidence that online retailers do not respond to local demand shocks, but offline retailers indeed respond to local shocks, even if they are indirectly affected.

We organize our paper as follows: Section 2 introduces our data and stylized facts on consumer heterogeneity. Section 3 leverages the Wayfair Decision to provide descriptive evidence of its impacts on sales, firm pricing, and heterogeneous substitution elasticities. Section 4 combines these new insights into a structural demand and supply model to evaluate welfare effects. Section 6 discusses the identification strategy and estimation results. We conduct our counterfactual exercises in Section 7 and conclude in Section 8.

2. Data

Our primary data set is the University of Chicago Kilts Center's archive of Numerator data, 2017-2021. Numerator is a marketing research firm that has surveyed two million U.S. households since 2017. Panelists record their offline purchases by uploading photos of their receipts; they also provide their email addresses so that Numerator can collect online receipts. We provide examples of original receipts in Appendix A.1. Each shopping receipt provides detailed information about the transaction, including the date, item descriptions, quantities, prices, and taxes. We also learn retailer identities and store addresses from these receipts. The Kilts Center archives all online receipts and a subset of offline receipts from purchases in the non-durable goods sectors. The data cover all major retailers in fast-moving consumer goods, such as Walmart, Target, and Amazon.com. Retailers specializing in durable goods (e.g., Apple Store, Best Buy, and IKEA) are excluded. We summarize the expenditures across all sectors and their online shares in Table A.2. Numerator also provides rich demographic information about the panelists, such as income level, car ownership, and ZIP. We demonstrate the representativeness of Numerator Data by comparing aggregate demographic statistics with US Census in Table A.3.

Using Numerator data has significant advantages. Numerator data offer omnichannel coverage, detailed prices, and tax information. We compare Numerator data with existing data sets like those of Nielsen, Comscore, Forrest, and Visa in Appendix A.2 and summarize the key differences in Table 1. We conclude that Numerator data are the best fit for our purposes.[4]

Table 1: Comparison of Numerator Data Set to Other Data Sets

	Omnichannel	Price	Tax	Retailer info
Numerator	✓	✓	✓	✓
Nielsen	✓	✓	✗	✗
Comscore	✗	✓	✓	✓
Forrest	✗	✗	✗	✗
Credit Card	✓	✗	✗	✓

To measure the level of sales tax change, we collect county-level sales tax rates from sales-tax.com. To understand the logistics costs borne by retailers, we use Infogroup data to locate the fulfillment centers of major retailers. We use this information in our supply estimation in Section 5. Appendix A.5 visualizes the distribution of the fulfillment centers of the retailers.

2.1 Sample Construction

In most of this paper, we restrict our attention to the pet-food retailer sector due to (1) data coverage and (2) its relevance to the Wayfair shock. First, although furniture and electronics retailers are the main subjects of the Supreme Court's case, we do not yet have full coverage of offline receipts from those durable sectors. Second, some non-durable sectors are not subject to the Wayfair shock. For instance, grocery food is nearly tax-free across all channels and states so the Wayfair Decision has no impact on the grocery industry.[5] We present the summary statistics of raw receipts containing at least one pet-food item in Table 2. Among households that have purchased at least one pet-food item, the average annual spending on pet food is $442 for the rich and $350 for the poor.

There are 7,104 unique retailers in the Numerator data. We identify the major players in the pet-food market in Table 3. We summarize three main takeaways. First, the market is highly

[4]We acknowledge some concurrent research also uses these new Numerator data. For example, Sullivan (2022) uses Numerator data to study the food-delivery industry. Song (2022) uses Numerator data to study the welfare effect of Supplemental Nutrition Assistance Program (SNAP) policies.

[5]We show this through a placebo test in Section 3.2.1.

Table 2: Summary Statistics of the Pet-Food Market (2017-2021)

	Mean	Sd	Min	25%	50%	75%	Max
Receipts (N = 26025201)							
by State-Month	8490	8783	178	2106	5641	11386	44483
by ZIP-Year	190	271	1	19	77	251	3370
Household (N = 1282137)							
by ZIP	38	58	1	3	11	47	977
Unique Store Visited (N = 120479)							
by ZIP	57	67	1	10	31	83	973

Note: Source: Numerator data (2017-2021). Receipts data: a collection of receipts containing at least one pet-food item; We conduct descriptive analysis on the State-Month level and estimate the structural model on the ZIP-Year level. Household data: headcounts of unique households that purchase pet food ; Unique store data: unique stores selling pet food based on receipts data.

concentrated, with the top five sellers accounting for 70% of the market. Second, the online market share is high. Two of the top three retailers are e-commerce retailers. The online stores of PetSmart and Petco are among the top 12 sellers. Third, the Wayfair Decision significantly affects the online market. The top three online sellers are all subject to the Wayfair shocks. We present the details of how each retailer is affected in Table 4.

2.2 Consumer Heterogeneity in the Online Market

We first document how online market shares differ by geographic region and demographic group. We group ZIP into 20 bins by the number of major offline stores available within a 50 km radius of each ZIP centroid. Figure 1a plots each bin's average online market share. We find that households located in rural areas, where fewer offline shopping stores are available, are more likely to shop online.

We then document online expenditure differences by demographic group, conditional on the offline (local) shopping environment. We run a linear probability model on shopping receipts. The dependent variable is an indicator of whether the receipt is an online receipt and the explanatory variables are a set of demographic characteristics, as in the following regression:

$$Pr(online\ receipts|X) = \Phi(\mathbf{X}\beta + F_{ZIP}). \tag{1}$$

X is a set of demographic dummies including income, education, gender, and car ownership;

Table 3: Top Pet-Food Retailers

#	banner	market share (%)	Online Store	Subject to Wayfair Decision
1	Walmart	24.7	✗	✗
2	Amazon.com	24.0	✓	✓
3	Chewy.com	8.0	✓	✓
4	Petco	6.6	✗	✗
5	PetSmart	6.5	✗	✗
6	Target	4.2	✗	✗
7	Costco	2.4	✗	✗
8	Sam's Club	2.2	✗	✗
9	Kroger	1.7	✗	✗
10	Pet Supplies Plus	1.7	✗	✗
11	Walmart.com	1.5	✓	✓
12	Meijer	1.4	✗	✗
13	Petco.com	0.6	✓	✗
14	Petsmart.com	0.5	✓	✗
15	All others	14	-	-

Note: We identify major retailers in the pet-food market. We treat online and offline stores belonging to the same retailer as a separate banners. Banner names without .com stand for offline stores. We treat sales from other retailers as an outside option. Source: Numerator.

and F_{ZIP} is the ZIP fixed effects. Figure 1b plots the coefficients of regression (1).

We find that conditional on ZIP, high-income households are 50% more likely to shop online than low-income households. We also find young households and highly educated households are more likely to shop online.[6] However, the consumer heterogeneity shown in Figure 1 does not directly translate into welfare measurement. To measure welfare, we need to know how substitutable online stores are for different groups of people. We document descriptive evidence of such substitution patterns in Section 3.2.3 and deploy a full structural estimation in Section 4.1.

3. Descriptive Evidence of Wayfair Decision

We first clarify the terminology used in our analysis in Figure 2. The primary entities of our study are retail firms and their stores. We use "firms" to refer to retailers rather than product manufacturers. Retailers source their products from various manufacturers and sell them in

[6] Sullivan (2022) finds similar patterns in the food-delivery industry.

Table 4: Main Online Retailers' Subjection to Wayfair Decision

Online Stores	Subjection	Coverage
Amazon.com	✗	✗
Amazon Marketplace	✓	Every state
Chewy.com	✓	Every state except TX, FL, MA, PA
Walmart.com	✗	✗
Walmart Marketplace	✓	Every state
Petco.com	✗	✗
Petsmart.com	✗	✗
Target.com	✗	✗

Note: Amazon.com, Walmart.com, Petco.com, Petsmart.com, and Target.com are not affected by the Wayfair Decision since they have fulfillment centers or physical stores in all states.

stores in three channels: (1) offline stores;[7]; (2) online stores, where the retailer owns the products and directly manages the prices (e.g., Amazon.com and Walmart.com refer to the online stores of Amazon and Walmart, over which the retailers have direct control); and (3) online marketplaces, where retailers provide a platform for transactions and can affect the final prices only indirectly (e.g., Amazon Marketplace refers to the platform provided for third-party sellers). In our setting, marketplaces exist only for Amazon and Walmart. We distinguish between online stores and marketplaces given the differential impacts the policy of interest has on them, as described in the next section.

3.1 Institutional Setting: The Supreme Court's Wayfair Decision

South Dakota v. Wayfair, Inc., Overstock.com, Inc., and Newegg, Inc. (hereafter, "Wayfair Decision") is a United States Supreme Court case decided on June 21, 2018. The decision gave states the choice of charging tax on purchases made from out-of-state sellers even if the seller does not have a physical presence in the taxing state. Before the decision, businesses without a physical presence in a particular state (i.e., those having no employees or property in said state) had no obligation to collect sales tax on out-of-state sales.[8] Out-of-state tax-free sellers typically include remote sellers such as Wayfair.com and online marketplace sellers. After the Wayfair

[7]We identify each store by its address (e.g., Walmart, 2300 Dixwell Ave, Hamden, CT 06514) and consider stores belonging to the same retailer as separate stores.

[8]The Supreme Court established this precedent in Quill Corp. v. North Dakota (usually known as the Quill Decision) in 1992.

Figure 1: Online Expenditure Heterogeneity

(a) Spatial Consumption Inequality

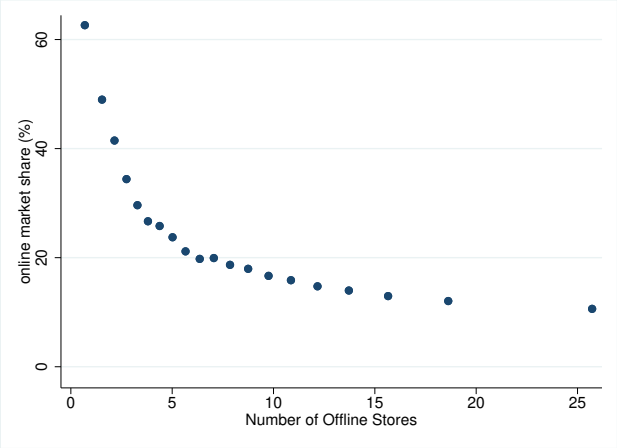

(b) Consumption Inequality by Demographics

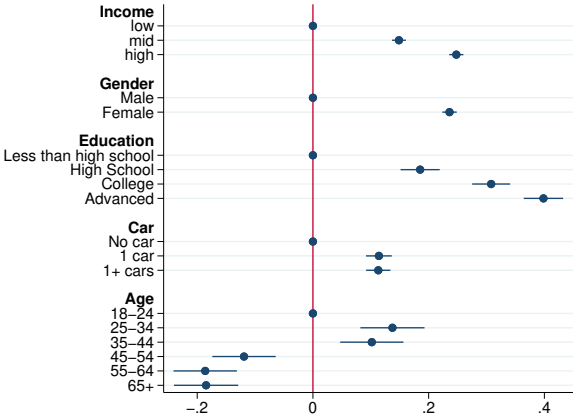

Note: (a) Binned scatter plot of online market shares versus the number of offline stores available within a 50 km radius of a ZIP centroid; (b) Coefficients plot of the linear probability regression of online shopping indicators on demographic variables and ZIP fixed effects, as in (1).

Decision, each state passed legislation to enforce sales taxes on remote and marketplace sellers.

We use an example to illustrate the policy change. An item, Rocco & Roxi dog treats, can be taxed differently across stores and time. If offline stores sell it, the local sales tax rate is always applied; if Amazon.com and Walmart.com sell it, it is likewise taxable since Amazon and Walmart have fulfillment centers or physical stores in every state. However, suppose Chewy.com sells it to a consumer who lives in Connecticut, where Chewy.com has no physical establishment. Such a sale would have been tax-free before the Wayfair Decision. It would have been similarly tax-free on Amazon or Walmart Marketplace since these marketplaces had no physical nexus before the Wayfair Decision. After the Wayfair Decision, Chewy.com and the Amazon and Walmart Marketplaces began collecting a tax of 6.35% based on the different tax-implementation timelines for remote and marketplace sellers set by the Connecticut legislature.

This drastic change in the online sales tax regime has had a large impact on the treated retailers[9]. In particular, it created rich variation in tax shocks at the firm-county-month level. First, the online sales tax change, even for the same firm, could be anywhere from 0% to 12.5%

[9]Appendix A.6 presents more narrative evidence of the significant effect of the Wayfair Decision.

Figure 2: Terminology

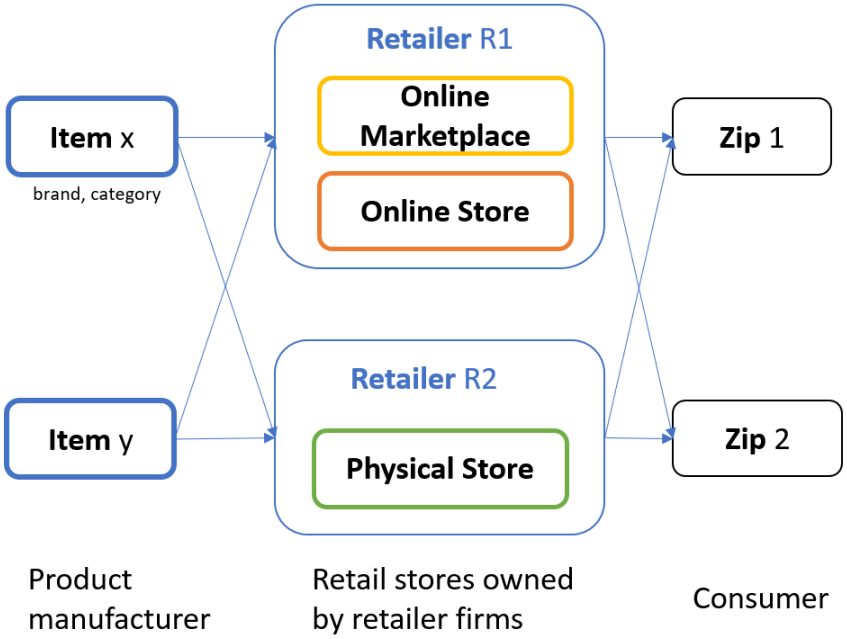

due to the variation in U.S. local sales tax rates, as shown in Figure 3.[10] Second, the timing of tax change for each state is between July 2018 and January 2023 since each state, according to its legislative process, set up separate effective dates for remote sellers and marketplace sellers. In Figure 4, we plot the staggered adoption date of each state against the state's average combined sales tax (state-level plus local-level sales tax). The graph helps to eliminate the concern that the adoption time may be endogenous. If that were the case, states with higher tax rates would be more likely to adopt tax changes in an earlier period. However, we do not observe those patterns in Figure 4.

We use the rich variation of sales tax change due to this institutional setting to identify the online demand. This identification strategy is formally established by Zoutman et al. (2018) and Dearing (2022). We discuss the identification details in Section 6.2.

Our empirical analysis focuses on the county-store level, and we construct monthly store-level sales and price index series by performing the following steps.

First, we treat Amazon(Walmart).com and Amazon(Walmart) Marketplace as different stores. For each transaction from Amazon or Walmart, we impute whether it comes from the market-

[10]Although there might be some local sales tax change after the Wayfair Decision, that would not affect our results since we are interested in the differential local sales taxes faced by online and offline sellers before the Wayfair Decision.

Figure 3: Sales Tax Distribution

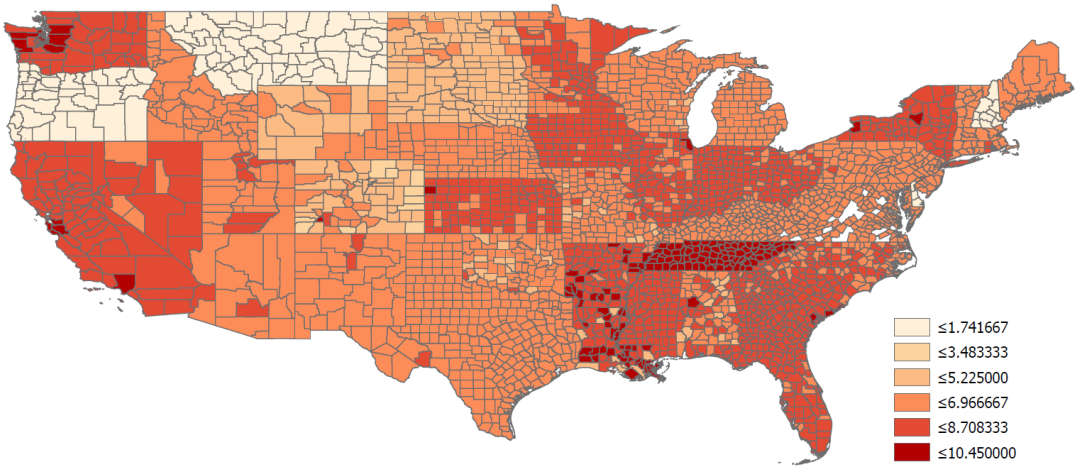

Note: County sales tax rates in the contiguous United States in 2018. Data collected from sales-tax.com.

place by looking at the item brand. If the item's brand is tax-free before the Wayfair Date, we classify this transaction as sales from the marketplace. We present the details of our marketplace imputation in Appendix A.7.

Second, some online sellers such as Chewy.com adjust their sales tax rates at a pace that does not perfectly adhere to the state-level deadline. Therefore, we identify the month each seller adjusts its sales tax in each state by implementing a tax-rate jump-detection algorithm using the sales tax information on receipts. We find that although some retailers deviate from the state deadlines in some states, most of the adjustments follow the state deadlines. We discuss the details of our tax-rate jump-detection algorithm in Appendix A.8.

Third, we calculate the store-level price as the unit price per ounce of pet food following Arcidiacono et al. (2020),

$$P_{s,t} = \frac{\sum_i p_{i,s,t} q_{i,s,t}}{\sum_i q_{i,s,t}}, \tag{2}$$

where each i stands for an item defined as a unique brand-category-unit combination; $p_{i,s,t}$ is the unit price per ounce; and $q_{i,s,t}$ is total ounces of item i sold in store s in month t. We discuss the details of our price index in Appendix B.1. We also count the unique number of items sold in store s in month t as $N_{s,t}$ to study retailers' potential adjustment on the assortment margin.

Figure 4: Staggered Adoption of the Wayfair Decision

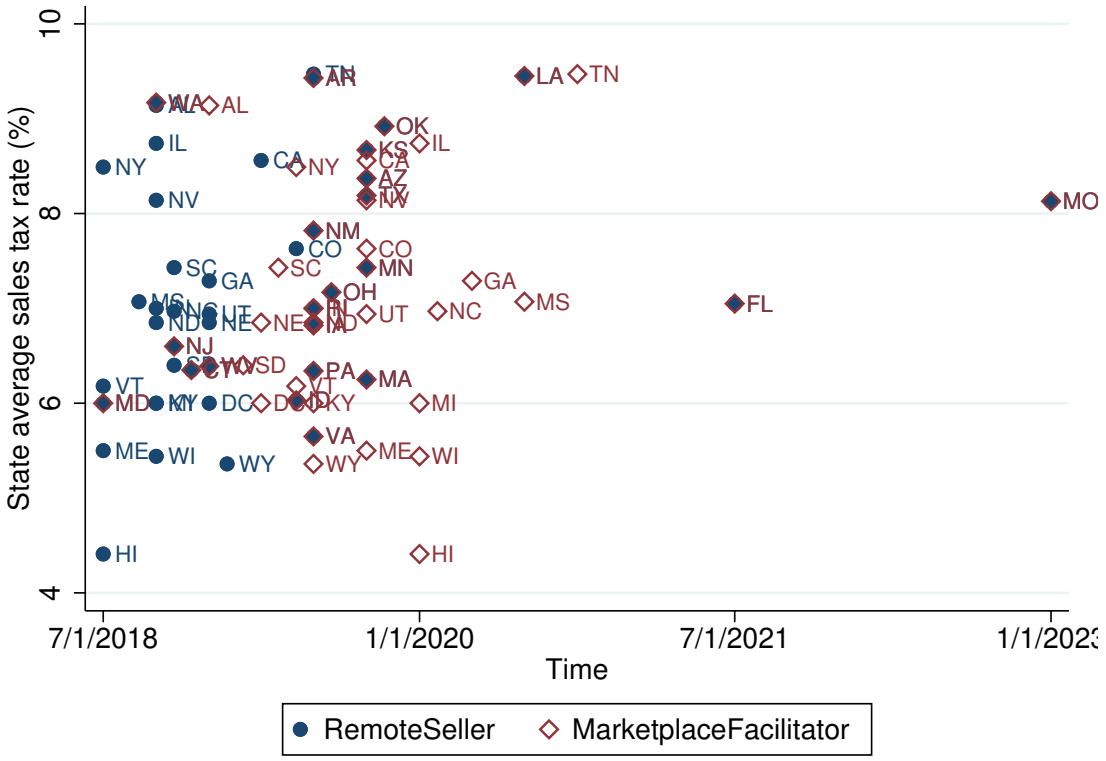

Note: Effective dates of sales tax implementation for remote sellers and marketplace sellers in each state. Data collected from state-level Internal Revenue Service (IRS) websites. Delaware, Montana, New Hampshire, and Oregon are omitted since they are tax-free states.

3.2 How Substitutable Are Online Stores?

In this section, we first deploy an event-study framework leveraging staggered adoption to visualize the effect of the Wayfair Decision on sales and firm responses. Section 3.2.1 documents the sales and pricing responses of the online sellers; Section 3.2.2 studies the sales and pricing responses of the indirectly treated sellers. Finally, in Section 3.2.3, we provide reduced-form evidence of the heterogeneous substitution elasticities among stores using both the timing and the level of sales tax changes.

3.2.1 Revenue and Pricing Responses of the Treated Retailers

In the following event study, we aggregate sales of firms on each channel to the state level and compare the trends of treated online sellers to indirectly treated ones using the following spec-

ification:

$$\ln Y_{f,c,r,t} = \sum_{k=-10}^{10} \beta_k \mathbb{1}_{f,c,r,k} + \delta_{f,c,r} + \delta_{f,c,t} + \delta_{r,t} + \epsilon_{f,c,r,t} \tag{3}$$

Here $Y_{f,c,r,t}$ are the outcome variables of interest. They are either the pre-tax sales $S_{f,c,r,t}$ or pre-tax average prices $P_{f,c,r,t}$ of firm f in channel $c \in$\{online store, online marketplace, offline\} made in state r, month t.[11] $\mathbb{1}_{f,c,r,k}$ is an indicator variable that takes 1 if there is a firm-channel sales tax change in month $t+k$. We bin event times ≥ 10 and ≤ -10 together. We also control for firm-channel-state fixed effects $\delta_{f,c,r}$, firm-channel-month fixed effects $\delta_{f,c,t}$, and state-month fixed effects $\delta_{r,t}$. We cluster standard errors at the state level.

Figure 5a plots the coefficients of β_k from estimating (3) with the log of pre-tax sales as dependent variables ($Y_{f,c,r,t} = S_{f,c,r,t}$). The figure has three implications. First, the Wayfair Decision significantly decreases the sales of the treated online sellers. On average, the Wayfair Decision decreases the sales revenue by 5% to 7%. Second, the almost-zero pre-trend coefficients suggest consumers do not react to the tax shock in advance. This is due to either the unexpected nature of the shocks or the non-hoarding behavior related to pet food. Lastly, there might exist a short-run salient effect of the tax shock. The sales do not immediately drop after the shock. It takes five months for the sales to stabilize.

Recent work in the applied econometrics literature suggests that an event study in ordinary least squares (OLS) might be biased when the treatment effects are heterogeneous (e.g., Borusyak et al. (2022), de Chaisemartin and D'Haultfoeuille (2022), Callaway and Sant'Anna (2021)). We adopt the robust and efficient imputation estimator from Borusyak et al. (2022) for our event study and present the results in Figure A.9. We arrive at similar results as those presented in Figure 5a.

To investigate the pricing response of the treated retailers, Figure 5b plots the coefficients of β_k from estimating (3) with the log of pre-tax prices as dependent variables ($Y_{f,c,r,t} = P_{f,c,r,t}$). We do not find online retailers specifically responding to local sales tax, partly because online retailers tend to set up uniform pricing nationwide, as confirmed by our store price discussion in Appendix A.8.1.

We also run a set of placebo tests to ensure that the revenue effects we find are not unique to the pet-food sector. Specifically, we run a regression of (3) on Amazon.com and Walmart.com

[11] In this section, we first treat the same online retailer in different states as separate online stores to investigate the possibility of online price discrimination by states. We calculate the online channel-state specific prices using (2).

Figure 5: Revenue and Pricing Responses of the Treated Retailers

(a) Log (pre-tax) Sales

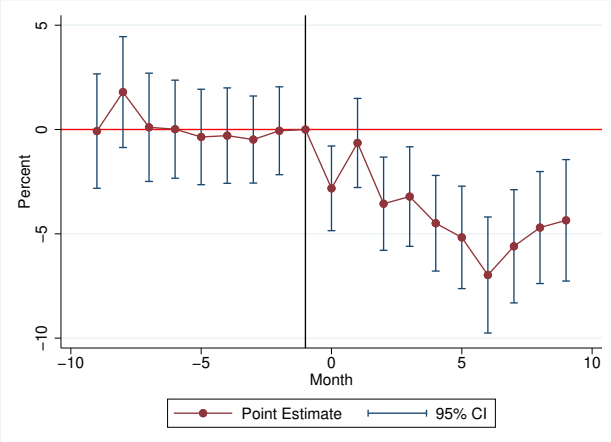

(b) Log (pre-tax) Prices

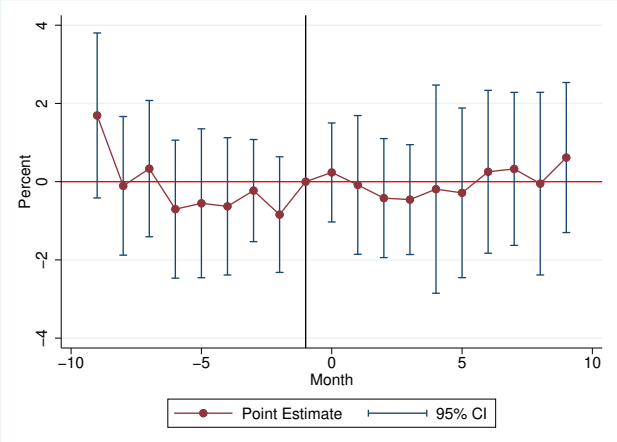

Note: Plot of β_k coefficients of our event-study regression (3). We regress pre-tax sales and pre-tax prices of firm-channel-state-month sales on the leads and lags of the tax shock, controlling for firm-channel-time, firm-channel-state, and state-time fixed effects.

over all sectors that are available in our data.[12] We find a significant negative impact on other taxable sectors such as health & beauty and home & garden and present these results in Figure A.10. Consistent with our expectations, we find no significant effect of the Wayfair Decision on non-taxable sectors such as grocery and baby. Figure A.11 illustrates these results. We report the event study of all other sectors in Appendix A.8.4.

3.2.2 Revenue and Pricing Responses of Indirectly Treated Retailers

In this section, we investigate the firm pricing and assortment responses of indirectly treated[13] retailers in each state by running a regression similar to (3):

$$\ln Y_{f,c,r,t} = \sum_{k=-10}^{10} \beta_k \mathbb{1}_{r,t,k} + \delta_{f,c,r} + \delta_{f,c,t} + \epsilon_{f,c,r,t}. \tag{4}$$

Here $Y_{f,c,r,t}$ is either pre-tax sales revenues $S_{f,c,r,t}$ or pre-tax average prices $P_{f,c,r,t}$ of indirectly treated firm f in channel $c \in$ {online store, online marketplace, offline} made in state r, month t.

[12] We do not include all remote sellers in this analysis because, as mentioned in Section 3.1, some remote sellers do not adjust their sales tax according to state-level deadlines.

[13] Indirectly treated retailers' sales tax do not change due to the Wayfair Decision. We consider these retailers indirectly treated because other retailers in the market were treated by the Wayfair Decision.

The indirectly treated sellers can be treated multiple times since remote sellers and marketplace sellers may adjust their sales tax rates at different time periods. To simplify the illustration, we choose $\mathbb{1}_{r,t,k}$ as the indicator of whether state r passes the remote seller adjustment deadline at month $t+k$. We drop the state-month fixed effects because they are collinear with the indicator function. We cluster standard errors at the state level.

Figure 6: Revenue and Pricing Responses of Indirectly Treated Retailers

(a) Log (pre-tax) Sales

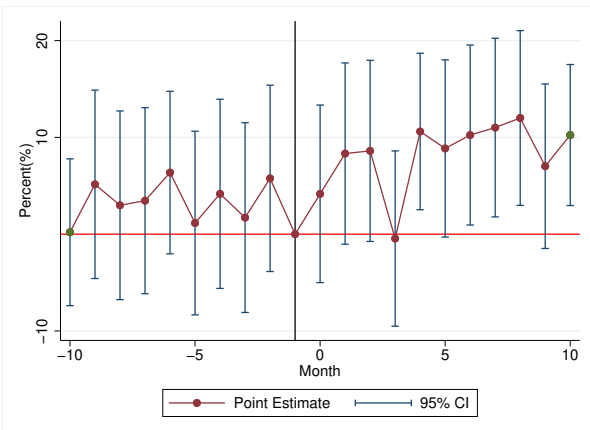

(b) Log (pre-tax) Prices

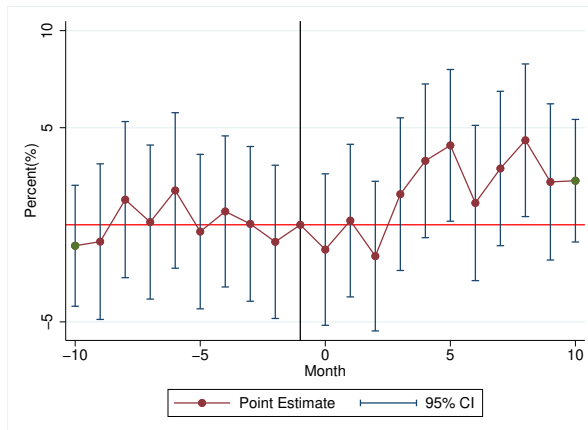

Note: Plot of β_k coefficients of our event-study regression on the log of pre-tax sales, as in (4). We regress firm-channel-state level pre-tax sales and pre-tax prices on the leads and lags of the tax shock, controlling for firm-channel-time and firm-channel-state fixed effects.

Figures 6a and 6b plot the coefficients of β_k from estimating (3) with the log of pre-tax revenues ($Y_{f,c,r,t} = S_{f,c,r,t}$) and log of pre-tax prices ($Y_{f,c,r,t} = P_{f,c,r,t}$) as dependent variables, respectively. After the Wayfair Decision, we find the indirectly treated retailers increase their pre-tax sales revenue by 10 % and increase their prices by 2%.

Of course, retailers could make adjustments other than prices, such as assortment changes. We address this concern by looking at the unique number of items sold in stores. We estimate (3) with the log of unique items sold as dependent variables ($Y_{f,c,r,t} = N_{f,c,r,t}$) and plot the coefficients of β_k in Figure A.14. We find retailers generally do not change their assortment in response to the Wayfair Decision.

3.2.3 Heterogeneous Substitution Elasticities

In this section, we provide reduced-form estimates of substitution elasticities among online stores and stores belonging to the same retailer using both the timing and the level of the sales tax change.

Cross-Channel Substitution

Figure 7: Heterogeneous Spillover 1: Online Versus Offline

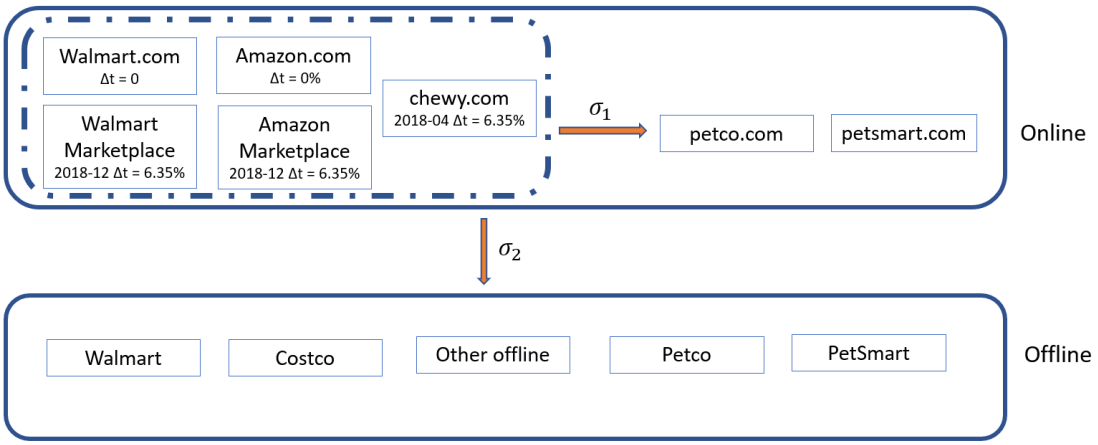

Note: Tax shocks directly affect sellers in the dashed box. When after-tax prices increase, consumers substitute away from treated online sellers into (1) indirectly treated online sellers and (2) indirectly treated offline sellers.

The differential exposure of the Wayfair Decision helps us learn about the substitution elasticities of online-online and online-offline. We group all retailers into three groups: treated online (including Amazon.com, Walmart.com, and Chewy.com),[14] indirectly treated online (including Petco.com and PetSmart.com), and indirectly treated offline (including all offline retailers). As the sales tax increases, we expect consumers to substitute from treated to indirectly treated groups, controlling for price changes. We run the following regression to estimate the

[14] In this section, we do not distinguish marketplace sellers from the marketplace owner and instead treat them as one seller. We do this because if we see consumers substitute from Amazon Marketplace into Amazon.com, we do not know whether this is due to their both being online stores or they both belonging to Amazon. We distinguish between these two in our structural estimation.

reduced-form cross-substitution elasticities:

$$\Delta \ln\left(\frac{S_{c,t}^{on,U}}{S_{c,t}^{on,T}}\right) = (1-\sigma_1)\,\Delta \ln\left(\frac{P_{c,t}^{on,U}}{P_{c,t}^{on,T}}\right) + \epsilon_{1,c,t} \tag{5}$$

$$\Delta \ln\left(\frac{S_{c,t}^{off,U}}{S_{c,t}^{on,T}}\right) = (1-\sigma_2)\,\Delta \ln\left(\frac{P_{c,t}^{off,U}}{P_{c,t}^{on,T}}\right) + \epsilon_{2,c,t}. \tag{6}$$

Here $S_{c,t}^{on,T}, S_{c,t}^{on,U}, S_{c,t}^{off,U}$ are the total after-tax sales of the treated, indirectly treated online, and indirectly treated offline stores, respectively, in county c, month t. $P_{c,t}^{on,T}, P_{c,t}^{on,U}, P_{c,t}^{off,U}$ are the after-tax group price indexes of the treated, indirectly treated online, and indirectly treated offline stores, respectively, in county c, month t. We calculate the change in the group price index as the sales-weighted change in the prices of each firm within the group (the Laspeyres price index)[15]. To estimate the reduced-form substitution elasticity, we regress the relative change of the after-tax sales on the relative change of the tax-included price index as in (5) and (6). $\sigma_1 > \sigma_2$ implies a higher substitution elasticity between online stores and vice versa. Since prices are endogenous, for example, if a store providing better service charges a higher price, the OLS estimates might be biased upwards. Hence, we use tax shocks as the instrument of relative prices. We choose a ten-month interval to calculate changes. We cluster standard errors at the county level.

Table (5) presents the estimation results of (5) and (6). Columns (1) and (3) report a positive substitution elasticity of the OLS estimates, which suggests the price endogeneity problem could be severe. We report the IV regression results in columns (2) and (4). As shown in the table, the substitution elasticities fall into the reasonable range. We find the online-online substitution elasticity is higher than the online-offline substitution elasticity.

Within-Firm Cannibalization

The Wayfair Decision also provides suggestive evidence of within-firm cannibalization. We ask whether, compared with other offline stores, Walmart offline stores particularly benefit from the tax shock on the Walmart Marketplace.[16] We divide offline sellers into Walmart and others

[15] We present the details of the construction of this price index in Appendix B.1

[16] Similarly, we do not distinguish between Walmart.com and Walmart Marketplace since the close substitution could be due to both online channels.

Table 5: Cross-Channel Estimation Results

	$\Delta_5 \ln(S^{on,U}_{c,t}/S^{on,T}_{c,t})$ OLS	$\Delta_5 \ln(S^{on,U}_{c,t}/S^{on,T}_{c,t})$ IV	$\Delta_5 \ln(S^{off,U}_{c,t}/S^{on,T}_{c,t})$ OLS	$\Delta_5 \ln(S^{off,U}_{c,t}/S^{on,T}_{c,t})$ IV
$\Delta_5 \ln(P^{on,U}_{c,t}/P^{on,T}_{c,t})$	0.238***	-2.483**		
$\widehat{1-\sigma_1}$	(0.0544)	(0.851)		
$\Delta_5 \ln(P^{off,u}_{c,t}/P^{on,T}_{c,t})$			0.282***	-0.693*
$\widehat{1-\sigma_2}$			(0.0331)	(0.324)
County FE	Yes	Yes	Yes	Yes
Observations	9013	9013	69850	69850

Note: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors in parentheses. Reduced-form cross-group substitution elasticities estimation. We regress the relative change in the price index on the relative change in the market shares. The first two columns report the online-online substitution elasticities. The last two columns report the offline-online substitution elasticities.

and estimate the following regression:

$$\Delta \ln\left(\frac{S^{W,off}_{c,t}}{S^{W,on}_{c,t}}\right) = (1-\sigma_3)\,\Delta \ln\left(\frac{P^{W,off}_{c,t}}{P^{W,on}_{c,t}}\right) + \epsilon_{r,t}, \tag{7}$$

$$\Delta \ln\left(\frac{S^{\text{-}W,off}_{c,t}}{S^{W,on}_{c,t}}\right) = (1-\sigma_4)\,\Delta \ln\left(\frac{P^{\text{-}W,off}_{c,t}}{P^{W,on}_{c,t}}\right) + \epsilon_{r,t}. \tag{8}$$

Here $S^{W,on}_{c,t}, S^{W,off}_{c,t}, S^{\text{-}W,off}_{c,t}$ are the total after-tax sales of Walmart online stores, Walmart offline stores, and other offline stores, respectively, in county c, month t. $P^{W,on}_{c,t}, P^{W,off}_{c,t}, P^{\text{-}W,off}_{c,t}$ are the after-tax price indexes of Walmart online stores, Walmart offline stores, and other offline stores, respectively, in county c, month t. Similarly, we calculate the group price index as the sales-weighted change in the prices of each firm within the group (the Laspeyres price index). To compare the substitution elasticity within a single retailer with that of stores belonging to different retailers, we regress relative after-tax sales on relative tax-inclusive prices, as in (7) and (8). $\sigma_3 > \sigma_4$ suggests a higher substitution elasticity among store chains and vice versa. We use tax shocks as instruments for relative prices to eliminate the endogeneity concerns of price settings.

We find that the substitution elasticity between Walmart offline stores and Walmart.com is higher than that between other stores and Walmart.com ($\sigma_3 > \sigma_4$). This suggests when Walmart Marketplace is taxed, consumers are more likely to switch to Walmart offline stores than other offline stores.

Figure 8: Heterogeneous Spillover 2: Cannibalization

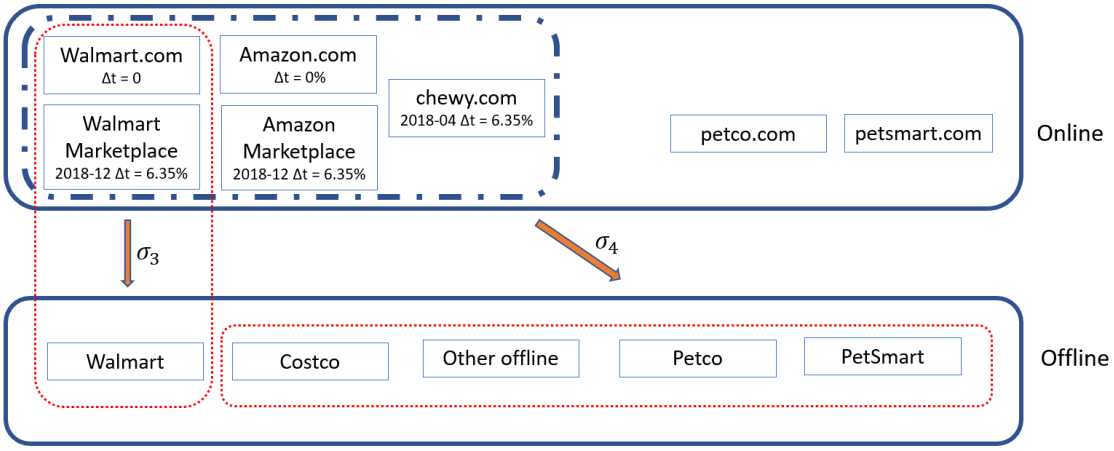

Note: Among offline sellers, Walmart may particularly benefit from Walmart Marketplace shocks.

Descriptive Evidence: Summary

The descriptive analysis in this section yields four takeaways: (1) the Wayfair Decision significantly decreases the sales of treated sellers; (2) online retailers set nearly uniform pre-tax pricing and do not respond to local tax changes; (3) indirectly treated offline retailers increase their prices by a small amount and do not adjust their assortment in response to the Wayfair Decision; and (4) the substitution elasticities among stores are heterogeneous: consumers substitute more across retailers if stores belong to the same online channel or retail firm.

However, these reduced-form estimates of substitution elasticities are unable to be rationalized by a complete demand system and therefore cannot be used to evaluate welfare directly. Hence, we develop a demand and supply model in the next section. The complete structural model incorporates the consumer heterogeneity in Section 2 and the flexible substitution we documented here and enables us to conduct counterfactual welfare analysis.

4. Demand and Supply in Pet-Food Retail Market

We now develop a demand and supply model for the spatial retailing industry with the presence of online shopping. We use the demand model to rationalize the different substitution elasticities we found in 3.2 and to estimate the heterogeneous preferences for online shopping among different demographic groups. The key substitution elasticities in the model allow us to bridge

Table 6: Within-Firm Substitution Estimation Results

	$\Delta_5 \ln(S^{on,u}_{c,t}/S^{on,T}_{c,t})$ OLS	$\Delta_5 \ln(S^{on,u}_{c,t}/S^{on,T}_{c,t})$ IV	$\Delta_5 \ln(S^{off,u}_{c,t}/S^{on,T}_{c,t})$ OLS	$\Delta_5 \ln(S^{off,u}_{c,t}/S^{on,T}_{c,t})$ IV
$\Delta_5 \ln(P^{on,u}_{c,t}/P^{on,T}_{c,t})$	-0.171	-18.87**		
$\widehat{1-\sigma_3}$	(0.0908)	(7.249)		
$\Delta_5 \ln(P^{off,u}_{c,t}/P^{on,T}_{c,t})$			-0.597***	-9.552*
$\widehat{1-\sigma_4}$			(0.0847)	(4.803)
County FE	Yes	Yes	Yes	Yes
Observations	32866	32866	25000	25000

Note: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors in parentheses. Reduced-form cannibalization estimation. We regress the relative change in the price index on the relative change in the market shares. The first two columns report the substitution elasticities between Walmart stores and Walmart online stores. The last two columns report the substitution elasticities between Walmart online stores to other stores.

the observed expenditure to a measure of consumer welfare. For the supply side, we model that firms incur different logistics costs to fulfill online and offline orders.

4.1 Demand Model

We define markets at the ZIP-year level indexed by (z,t). Within each market, each consumer i belonging to demographic group d chooses a pet store and then spends their budget there. The stores available to each consumer in each market are denoted by $j \in \mathbb{S}(z)$. Stores sell on different channels $C(s) \in \{\text{online,offline}\}$ and belong to different firms (retailers) $F(s) \in \mathbb{F}$. The indirect utility of choosing store j depends on after-tax prices, distance, quality, and idiosyncratic taste, which, suppressing the month subscript t, we specify as

$$V_{i,j,z} = \delta_{d,j,z} + \mu_{i,j} + \epsilon_{i,j,z}, \tag{9}$$

where $\ln\delta$ is the mean utility, common to everyone in the same location and demographic group; while μ and ϵ are the unobserved heterogeneous preferences. We specify the mean utility as follows:

$$\ln\delta_{d,j,z} = \ln y_d + A_{d,j} - \alpha_d \ln[p_j(1+\tau_{j,z})] - \beta_d \ln d_{j,z} + \xi_{d,j,z}. \tag{10}$$

Figure 9: Demand: Mean Utility

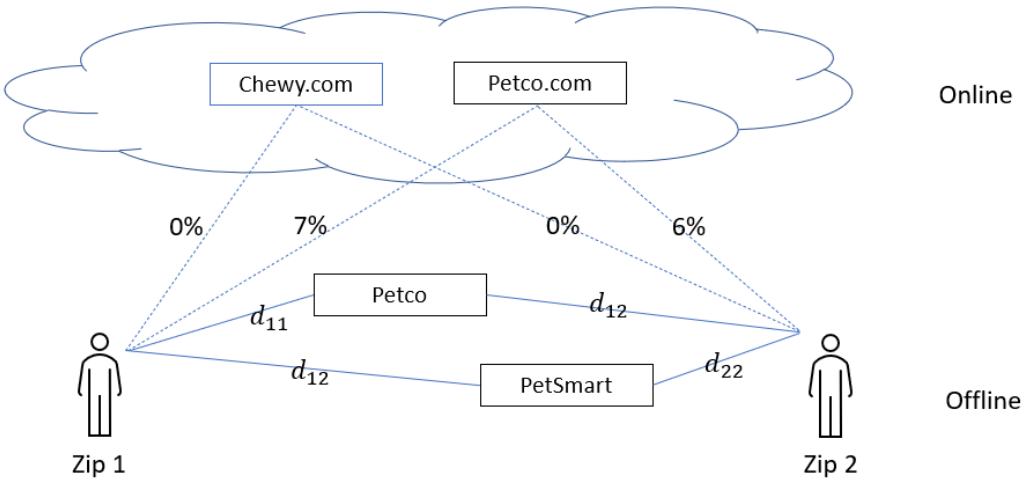

Note: Mean utility depends on store amenities, prices, taxes, and distance.

We now suppress the demographic subscript d to simplify our discussion. In the equation above, y is the budget for pet food,[17] A_j captures the store j quality: for online stores, it captures online amenities such as ease of search and comparison; for offline stores, it captures store-level characteristics such as proximity to shopping plaza, in-person service, and assortment. p_j is the measured store price index. $\tau_{j,z}$ is the sales tax rate: for offline stores, $\tau_{j,z} = \tau_j$ since sales tax only depends on the location of the stores; however, for online stores, $\tau_{j,z}$ is store-ZIP specific since online sales tax depends on the retailer's physical presence in a given state, the local tax rate, and the timing of a change in the sales tax rate due to the Wayfair Decision. We use log-logit specification because it can nest the canonical constant elasticity of substitution (CES) price index as a critical special case and separate prices and taxes so that we can estimate the price elasticity α only using tax variations, which we discuss in Section 6.2. $d_{s,z}$ describes the distance to the store: it is equal to a constant for online stores since all online stores are equally accessible to all customers. $\xi_{j,z}$ is the unobserved ZIP demand shifter.

[17] We explicitly address non-homothetic preference by specifying the discrete types of demographic groups that differ in their preferences. Consequently, budget and prices are separable within each group.

Figure 10: Demand: Idiosyncratic Preference

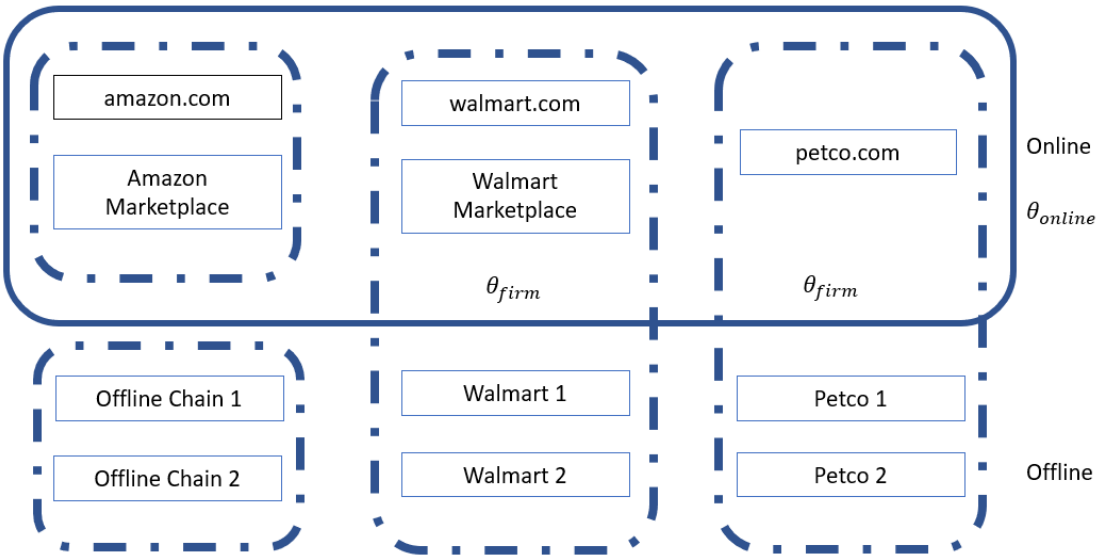

Note: Correlated idiosyncratic shocks introduce close substitution among channels and firms.

We introduce flexible substitution through $\mu_{i,j}$:

$$\mu_{i,j} = \epsilon_{i,F(j)} + \epsilon_{i,\text{on}} \mathbb{1}_{C(j)=\text{online}}. \tag{11}$$

$\epsilon_{i,F(j)}$ is the random firm preference shock and $\epsilon_{i,C(j)}$ is the random online preference shock. $\epsilon_{i,j,z}$ is the store idiosyncratic preference shock. The random online (firm) preference shocks will generate correlations within online channels (firms), which in turn introduce close substitution patterns among the stores belonging to the same channel (firm), as shown in Figure 10. Instead of assuming that random coefficients follow normal distributions, we specify the joint distribution of $\mu_{i,j} + \epsilon_{i,j,z}$ following the GEV distribution as follows:

Assumption 1 *Assume that $\epsilon_j \equiv \mu_{i,j} + \epsilon_{i,j,z}$ follows the GEV distribution with c.d.f:*

$$F(\epsilon_1, \ldots, \epsilon_N) = \exp\left(-\left(\sum_{j \in on} \lambda \mathbb{1}_{C(j)=on} e^{-\epsilon_j \theta_{on}} \right)^{\frac{1}{\theta_{on}}} - \sum_{f \in F} \left(\sum_{f} \left(1 - \lambda \mathbb{1}_{C(j)=on}\right) e^{-\epsilon_j \theta_f} \right)^{\frac{1}{\theta_f}} \right)$$

.

We use this GEV specification for the following reasons. First, it introduces closed-form mar-

ket shares and price-index expressions, whose values and derivatives are easy to compute.[18] Second, it generates a similar correlation structure as the random coefficients models as discussed in Section 4.2. Third, it succinctly captures the substitution patterns into easy-to-interpret nesting parameters θ_{on}, θ_f and λ. Under Assumption 1, the close substitution of stores in the same channel and in the same firm is captured by θ_{on} and θ_f, respectively. Here each online choice is split into two parts and placed separately into a channel nest and a firm nest. The parameter λ describes the weight of store presence in each nest. We then show that Assumption 1 gives the closed-form market share and expected maximized utility.

Proposition 1 *Under Assumption 1, the (after-tax) sales market share takes the form of*

$$S_{j,z} = \frac{\left(\lambda \mathbb{1}_{C(j)=on}\delta_{j,z}\right)^{\theta_{on}} \Phi_{on}^{1/\theta_{on}-1} + \left(\left(1-\lambda \mathbb{1}_{C(j)=on}\right)\delta_{j,z}\right)^{\theta_f} \Phi_f^{1/\theta_f-1}}{\Phi_z}, \tag{12}$$

where $\delta_{j,z}$ is the mean utility in (10). Φ_{on}, Φ_f are the inclusive values of the online channel and firm f nest, respectively: $\Phi_{on} \equiv \sum_{C(j)=on}(\lambda \mathbb{1}_{C(j)=\text{on}}\delta_{j,z})^{\theta_{on}}$, $\Phi_f \equiv \sum_{F(j)=f}(1-\lambda \mathbb{1}_{C(j)=\text{on}})\delta_{j,z})^{\theta_f}$. The denominator Φ is the sum of the inclusive values of all channel and firm nests in that market. $\Phi_z \equiv \Phi_{on}^{1/\theta_{on}} + \sum_f \Phi_f^{1/\theta_f}$. The GEV assumptions also give us the market-specific expected-max utility in a closed form,

$$W_z = \ln \Phi_z = \left(\left(\sum_{C(j)=on}\left(\lambda \mathbb{1}_{C(j)=\text{on}}\delta_{j,z}\right)^{\theta_{on}}\right)^{1/\theta_{on}} + \sum_f \left(\sum_{F(j)=f}\left(1-\lambda \mathbb{1}_{C(j)=\text{on}}\right)\delta_{j,z})^{\theta_f}\right)^{1/\theta_f}\right), \tag{13}$$

which is a nonlinear combination of all the mean utilities (including quality, prices, taxes, and distance) of the stores in ZIP z. These features closely follow McFadden (1980) and Feenstra (1995). We formally prove these results in Appendix B.2.

If prices are fixed, knowing the structural parameters $\alpha, \beta, \theta_c, \theta_f, \lambda$, and the observed market shares is sufficient to conduct a counterfactual analysis of consumer welfare. However, since firms endogenously set prices, we need to model supply in Section 5.

[18] In practice, we adopt automatic differentiation (AD) to evaluate the exact Jacobian function. Compared with alternative methods such as symbolic differentiation and numerical differentiation, AD has three main advantages: (1) it can be adjusted for any variations of choice sets; (2) it is free from numerical simulation errors; and (3) it is fast and efficient to compute.

4.2 Demand Model Discussion

Before introducing the supply model, we discuss the relationship between our demand system and CES, nested-CES, random coefficients models, and demand systems with local complementarity.

First, a critical special case of our demand system is CES. If there is no close substitution in either the online nest or the firm nests, $\theta_{on} = \theta_f = 1$, then our demand system will collapse into the canonical CES demand system. Following insights from Anderson and van Wincoop (2003) and Fally (2015), we can directly estimate the price elasticities and distance elasticities using the gravity equation estimation with two-way fixed effects as follows:

$$\ln S_{j,z} = F_j + F_z - \alpha \ln(1 + \tau_{j,z}) - \beta \ln d_{j,z} + \epsilon_{j,z}.$$

However, if rich substitution patterns exist beyond CES, we cannot estimate the elasticities using the previous equation.

Second, our cross-nested demand system is a generalization of the nested-CES demand models, which usually put strong priors on the nesting structure of each choice. For example, if $\lambda \to 1$, it reduces to the channel-nested CES demand, which assumes that only online stores have closer substitution. If $\lambda \to 0$, it reduces to the firm-nested CES model, which assumes that only stores belonging to the same firm have closer substitution. In our generalized cross-nested demand system, we do not impose such priors. The nesting structure can be overlapped and captured by λ, which we use the data to reveal.

Third, to generate a correlation matrix of all choices, our demand system is quantitatively similar to the random coefficients models, as in Berry (1994) and Berry et al. (1995). In Table 7, we compare the covariance matrix of our cross-nested demand model with that of the random coefficients model (with random coefficients on online channel dummies and firm dummies). The cross-nested model has the flexibility to match the covariance matrix with the random coefficients model. The model likewise remains a closed-form market share solution, saving us from numerical simulations. However, we acknowledge that the cross-nested model can handle arbitrary substitutions based only on categorical variables, not continuous variables. Therefore, it cannot nest the models with random coefficients on prices and distance in our setting.

Lastly, our demand curves derived as in (12) can easily accommodate local complementarity if researchers are willing to assume $\theta_{on}, \theta_f < 1$. However, this specification cannot be

Table 7: Cross-Nested and Random Coefficients Demand Systems

$COV(\epsilon_i, \epsilon_j)$	Random Coefficient	Cross-Nested (Approx.)
different channel, different firms	0	0
same online channel, different firms	σ_{on}^2	$\lambda^2(1-\frac{1}{\theta_{on}^2})$
different channel, same firms	σ_f^2	$(1-\lambda)(1-\frac{1}{\theta_f^2})$
same online channel, same firms	$\rho\sigma_{on}\sigma_f$	$\lambda^2(1-\frac{1}{\theta_{on}^2})+(1-\lambda)^2(1-\frac{1}{\theta_f^2})$

Note: The covariance matrix derived for the cross-nested model is an empirical approximation by Papola (2004). Cross-nested demand can be flexible to match the covariance matrix of the random coefficients model based on categorical variables but remain a closed-form market share solution.

micro-founded by previous random utility discrete-choice models.[19] In practice, we take the demand curves in (12) to the data and do not restrict the value of nesting parameters in our estimation. Our estimates suggest all $\theta_{\cdot} > 1$, which is consistent with our random utility model specifications.

5. Supply Model

In the supply model, we include retailers' fulfillment centers to learn about the logistics costs difference when fulfilling online and offline orders, as is shown in Figure 11. For offline sales, firms need to ship products from their fulfillment centers to their physical stores, and it is usually done by a full truckload (FTL); for online sales, firms need to ship products from their fulfillment centers to customers' addresses in parcels, and it is usually done by a specialty carrier like UPS. We do not directly observe shipping cost from our data sets. Instead, we infer shipping costs by inverting the equilibrium prices conditional on the shipping distance and estimated demand. The intuition is that conditional on consumer demand and the shipping distance, if a retailer sets higher prices for its online stores than its offline stores, then online logistics costs are higher. We discuss the details of this identification strategy in Section 6.2.

Assumption 2 *We assume the final cost at stores is a combination of the cost of goods from the*

[19] Global substitutes are a common feature of random utility discrete-choice models, apart from those that have special bundling assumptions, as in Gentzkow (2007). See Feng et al. (2018) for a more formal discussion of this issue.

Figure 11: Supply Model

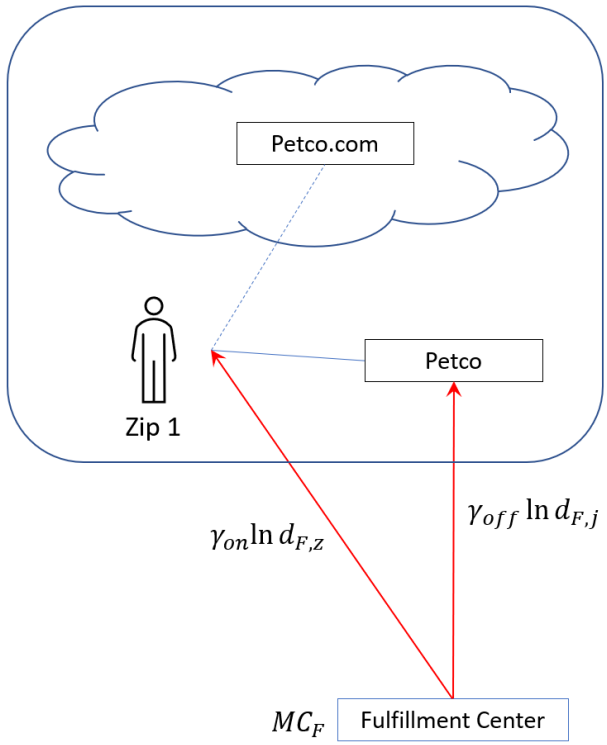

Note: Final costs = Marginal cost at the fulfillment center + logistics costs. The logistics costs differ by retail channel mode: while offline sales need to remit the logistics costs from the fulfillment center to the physical stores, online sales need to be shipped directly to a given household's address.

nearest fulfillment center and the logistics costs from the nearest fulfillment center:

$$\ln c_{f,s}^{offline} = \ln c_f + \gamma_{offline} \ln d_{f,s} + \eta_{f,r}^{offline}, \tag{14}$$

$$\ln c_{f,z}^{online} = \ln c_f + \gamma_{online} \ln d_{f,z} + \eta_{f}^{online}. \tag{15}$$

Here c_f is the marginal cost at all fulfillment centers belonging to firm f. For offline sales, $d_{f,s}$ is the distance from store s to the nearest fulfillment center and the logistics cost parameter is $\gamma_{offline}$. For online sales, $d_{f,z}$ is the distance from the consumer's ZIP centroid to the nearest fulfillment center and the logistics cost parameter is γ_{online}.

Recent literature suggests retailers set zone prices rather than store-specific prices (DellaVigna and Gentzkow (2019), Adams and Williams (2019), Butters et al. (2022)). To accommodate this concern, we make the following pricing-zone assumptions:

Assumption 3 *We assume each retailer sets uniform offline prices at the state level and sets a uniform nationwide price for online stores.*

Therefore, the supply-side structural errors in the aforementioned equations only show up at the pricing-zone level as $\eta_{f,r}^{offline}, \eta_{f}^{online}$. Given the cost structure and pricing-zone assumptions, we make the following assumption about firm pricing behavior:

Assumption 4 *We assume retail firms are engaged in the Bertrand-Nash pricing game: each firm sets up its offline and online zone prices to maximize its profit in each pricing zone. That is,*

$$\max_{\{p_{f,r}^{offline}\}, p_f^{online}} \pi_f = \sum_{z\in r}\sum_{s\in f} \frac{Y_z S_{s,z}(p_{f,r}^{offline}; p_{-1}, \tau)}{1+\tau_s}(1-\frac{c_{f,s}^{offline}}{p_{f,r}^{offline}}) + \sum_z \frac{Y_z S_{s,z}(p_f^{online}; p_{-1}, \tau)}{1+\tau_z}(1-\frac{c_{f,z}^{online}}{p_f^{online}}). \tag{16}$$

The first line is the sales profit from offline stores and the second line is the total sales from online stores. We believe Assumptions 2-4 are a reasonable approximation of the actual retail industry.

5.1 Supply Model Discussion

The supply-side assumptions allow us to infer firms' marginal cost at a fulfillment center, conditional on demand, distance, and observed prices. To simplify our illustration, we start with a pure offline seller's first-order condition (FOC) of (16):

$$\sum_{z\in r}\sum_{s\in f} \frac{Y_z}{1+\tau_s}\left(\frac{dS_{s,z}}{dp_{f,r}^{offline}}(1-\frac{c_{f,s}^{offline}}{p_{f,r}^{offline}}) + \frac{S_{s,z}c_{f,s}^{offline}}{\left(p_{f,r}^{offline}\right)^2}\right) = 0. \tag{17}$$

Thus we can solve for the marginal cost at the fulfillment center as

$$\frac{c_f \eta_{f,r}^{offline}}{p_{f,r}^{offline}} = -\frac{\sum_s\sum_z w_{s,z}\Delta_{f,c}^{s,z}}{\sum_s\sum_z w_{s,z} d_{c,s,z}^{\gamma_c}(1-\Delta_{f,c}^{s,c})}, \tag{18}$$

where $w_{s,z} = \frac{Y_z S_{s,z}/(1+\tau_{s,z})}{\sum_s\sum_z Y_z S_{s,z}/(1+\tau_{s,z})}$ are the tax-adjusted sales weights that we directly observe from the data; and $\Delta_{f,c}^{s,c} \equiv \frac{\alpha d \ln S_{s,z}}{d \ln \delta_{f,s}}$ is the transformed Jacobian of demand in (12). For an omnichannel

retailer, Δ also encapsulates the interactions between online and offline stores: when a firm increases its online prices, the demand for its offline stores might increase. Due to the closed-form solution, Δ is straightforward to compute, which helps us to gain computing efficiency in the estimation and counterfactual.

The supply model completes the industrial equilibrium. The general equilibrium is given taxes, demand and supply shifters, and firm conduct, the consumers solve utility-maximization problems which yields demand curves in (12) and firms solve profit maximization problems in (16) such that market is clear. The demand and supply shifters are unobserved to econometricians, which raises endogeneity concerns in estimation. We thus discuss our identification strategy in the next section.

6. Identification and Estimation

In this section, we lay out the identification and estimation procedures for our demand and supply model. Section 6.1 discusses the challenges of demand estimation for the online market and threats to our tax identification strategy. Section 6.2 establishes our identification strategy for demand and supply. Section 6.3 introduces our full estimation procedure. We present baseline estimation results in Section 6.4 and discuss robustness checks in Section 6.5.

6.1 Challenges of Online Demand Identification

Since prices are endogenous, we need to find demand instruments that credibly identify the price elasticity. One popular choice for price instruments is the so-called Hausman instruments (as in Hausman et al. (1994), Nevo (2001), DellaVigna and Gentzkow (2019), Handbury (2021)). The researchers use "same chain-other city" prices as instruments, assuming variations in prices reflect only the cost structure and are orthogonal to demand shocks. This exclusion restriction is hard to check since demand shocks may be correlated. Additionally, as we have shown in Figure A.8.1, online prices are mostly uniform, which may lead to weak instrument issues.[20] Therefore, we do not use Hausman instruments for our estimation.

In our setting, using local sales tax variations is particularly appropriate because there is rich variation at the firm-county-month level. We thus choose this method, treating the level and timing of tax adjustments as orthogonal to demand shocks. Yet this identification strategy will

[20] In Appendix A.10, we discuss this issue using the language of directed acyclic graphs (DAG).

fail if demand shocks are correlated with tax shocks. For example, the Wayfair Decision could "cause" consumers in a given state to increase their online demand and further delay state-level adjustment through the legislative process. Figure 4 relieves this concern by showing there is no such correlation between tax rates and adoption dates. Another threat to our identification strategy is that we assume that the tax elasticity is equal to the price elasticity. There exists some evidence that consumer behavior might not be fully salient to taxes (e.g, Chetty et al. (2009), Kroft et al. (2020)). We mention this issue in Section 6.5 but do not directly address it in this paper.

6.2 Identification

In this section, we formally establish our identification assumptions for the demand and supply estimation.

Demand: Price and Distance Elasticities

We first describe how we identify the price elasticity α_d and distance elasticity β_d, conditional on nesting parameters $\Theta = (\theta_{on}, \theta_f, \lambda)$. Denote geographical market (ZIP code) as z, time as t, and store (Amazon.com) as j. Consider the following inverted demand system:

$$\delta_{d,j,z,t}(\boldsymbol{p}, \boldsymbol{q}, \Theta) - \delta_{d,0,z,t}(\boldsymbol{p}, \boldsymbol{q}, \Theta) = A_{d,j,t} - \alpha_d g(p_{j,t}(1+\tau_{j,z,t})) - \beta_d h(d_{j,z}) + \xi_{d,j,z,t}. \tag{19}$$

In this system, $\delta_{d,j,z,t}$ is the inverted demand for demographic group $d \in$ {rich, poor} that depends on prices, quantity, and nonlinear demand parameters. g, h are known functions. This demand system nests logit, CES, random coefficients, and nested versions of them, as well as GNL, and GEV models.[21] For our specification, $g(x) = h(x) = \ln x$. Apart from prices, there may be concerns that local tax rates are correlated with local unobserved demand. For example, urban areas may have both higher unobserved demand for online shopping and higher tax rates. We alleviate this concern by including county-fixed effects. There may also be concerns that distance is correlated with unobserved demand. For example, firms might open more stores in regions with higher unobserved demand, thereby causing the distance to be correlated with unobserved demand. Those concerns are alleviated by our inclusion of store-time fixed effects.

[21] For example, in logit demand systems, $\delta_{d,j,z,t} = \ln s_{d,j,z,t}$ and $g(x) = x$. In CES demand systems, $\delta_{d,j,z,t} = \ln s_{d,j,z,t}$ and $g(x) = \ln x$.

(19) thus becomes

$$\delta_{d,j,z,t}(\boldsymbol{p},\boldsymbol{q},\Theta)-\delta_{d,0,z,t}(\boldsymbol{p},\boldsymbol{q},\Theta)=F_{d,j,t}+F_{c(z)}-\alpha_d g(p_{j,t}(1+\tau_{j,c(z),t}))-\beta_d h(d_{j,z})+\tilde{\xi}_{d,j,z,t}, \quad (20)$$

where $\tilde{\xi}_{d,j,z,t}$ is the demand residual after controlling for unobserved demand at the county and store-time levels. Proposition 2 (Dearing (2022)) formally states our identification strategy for price elasticity α:

Proposition 2 *If (i) $\mathbb{E}(\tilde{\xi}_{j,m,t}|\tau_{j,m,t})=0$, (ii) $\tau_{j,m,t}$ only affects demand through after-tax prices, and (iii) $\mathbb{COV}(g(\tilde{p}_{j,t}),\tau_{j,m,t})$ has full rank, then α is identified.*

Here Assumption (i) is the standard IV exclusion restriction. Assumption (ii) is the standard Ramsey restriction that says tax only enters the demand system through after-tax prices. If tax affects other variables in (19), such as quality A, then we cannot identify α without knowing the exact relationship between tax and quality. Assumption (iii) says we need enough variation in taxes and after-tax prices. It highlights the need for changes in the tax rates. If there is no tax shock, the tax-level variations will be absorbed in the county's fixed effects.

Our identification of distance elasticity β requires that this demand residual is independent of distance. Because of finite sampling, some stores belonging to the consumer choice set may have zero market share. Therefore, the inversion of (20) might not exist, which would lead to underestimating β due to sample selection. In Section 6.5, we address this issue using the Poisson pseudo maximum likelihood (PPML) estimator that is commonly used in the trade and spatial literature.

Demand: Nesting Parameters

To identify nesting parameters $\Theta=(\theta_{on},\theta_f,\lambda)$, we construct instruments in each market exploring the spatial variations of consumers' choice sets (usually referred as "BLP-style instruments"). They are the number of rival online stores $N^{rival}_{on,j,z}$, the number of non-rival online stores $N^{nonrival}_{on,j,z}$, and the number of same-firm stores $N_{f,j,z}$. We discuss the identification starting with the intuition from the nested-logit model (Berry (1994)).

First, if $\lambda=0$ or 1, then the GEV demand system reduces to a nested-logit demand system that can be directly estimated using the following regression:

$$\ln S_{j,z,t}-\ln S_{0,z,t}=F_{j,t}+F_{c(z)}-\alpha\ln(1+\tau_{j,z,t})-\beta\ln d_{j,z}+(1-\frac{1}{\theta_g})\ln S_{j|g}+\tilde{\xi}_{j,z,t}. \quad (21)$$

If $\lambda = 0$, then (21) corresponds to a firm-nested model, where $\theta_g = \theta_f$ and $\ln S_{j|g}$ becomes the within-firm market share. Since within-market shares are endogenous, one typical instrument is the number of stores belonging to the same firm, $N_{f,z}$. For example, if one market has more Walmart stores, then a given Walmart store will have a smaller within-firm market share if Walmart stores are close substitutes. Similarly, if $\lambda = 1$, then (21) becomes a channel-nested model, where $\theta_g = \theta_c$ and $\ln S_{j|g}$ becomes the within-channel market share. One typical instrument for $\ln S_{j|g}$ is the number of online stores $N_{on,z} = N^{rival}_{on,j,z} + N^{nonrival}_{on,j,z}$. For example, if one market has more online stores than another, then a typical online store will have a smaller within-channel market share if the online stores are close substitutes. In a more general case where λ is in between 0 and 1, there is no regression-like equation similar to (21). In this case, within-firm share and within-channel share are driven by the number of rival and non-rival stores simultaneously. Therefore, $N^{rival}_{on,j,z}$ and $N^{nonrival}_{on,j,z}$ can be used to pin down the splitting parameters λ.

We assume here the variation in rival online stores $N^{rival}_{on,j,z}$ comes from consumer awareness. For example, if one region does not have any offline Petco stores, Petco.com may not exist in consumer choice sets. However, one may ask whether this variation comes from measurement error because of finite sampling. We address this concern by proposing another BLP instrument in Section 6.5, leveraging the sales tax change of competitors, which is not subject to measurement error concerns.

Supply-side Parameters

We now turn to supply-side parameters $\Gamma = (\gamma_{online}, \gamma_{offline})$. The basic idea is that given the demand-side parameters and observed equilibrium, we can invert the firm's optimal pricing as in (17) to calculate the marginal cost at the firm's fulfillment centers based on our assumptions about firm conduct. Given a guess of Γ, we can compute the supply shifters $\widehat{\eta^{offline}_{f,d}}(\Gamma), \widehat{\eta^{online}_{f}}(\Gamma)$ according to (18).

$\gamma_{offline}$ is chosen such that conditional on demand and distance to fulfillment centers, marginal costs at all fulfillment centers of the same firm are as close as possible. The corresponding moment condition thus is

$$g^S_1(\gamma_{offline}) = \frac{1}{N_{fc}} \widehat{\eta^{offline}_{f,c}}(\gamma_{offline}). \tag{22}$$

Similarly, γ_{online} is chosen to match the inverted online marginal cost to the offline marginal

cost at the fulfillment centers. The moment condition is

$$g_2^S(\gamma_{online}) = \frac{1}{N_f}\widehat{\eta_f^{online}}(\gamma_{offline}). \tag{23}$$

6.3 Full Estimation Procedure

The full estimation procedure follows the nested fixed-point algorithm proposed in Berry et al. (1995). We first estimate demand-side parameters. Given a guess of nesting parameters $\Theta = (\theta_{on}, \theta_f, \lambda)$, we begin by obtaining the mean utility $\delta_{d,j,z,t}$ that solves the nonlinear equations in (12). We then run the following regression to obtain the unobserved demand residual $\tilde{\xi}_{d,j,z,t}$:

$$\ln \delta_{d,j,z,t}(\Theta) = F_{d,j,t} + F_c - \alpha_d \ln(1 + \tau_{j,c(z),t}) - \beta_d \ln d_{j,z} + \tilde{\xi}_{d,j,z,t}(\Theta). \tag{24}$$

We then use BLP-instruments $Z_{j,z,t} = (N_{on,j,z}^{rival}, N_{on,j,z}^{nonrival}, N_{f,j,z})$ to construct demand-side moment conditions: $\mathbf{g}^D(\Theta) = \frac{1}{N}\sum \tilde{\xi}_{j,z,t}(\Theta) Z_{j,z,t}$. We choose Θ to minimize the GMM objective function $\mathbf{g}^D(\Theta)' W \mathbf{g}^D(\Theta)$.

We subsequently estimate the supply-side parameters. Given the estimated demand-side parameters and a guess of supply-side parameters $\Gamma = \left(\gamma_{online}, \gamma_{offline}\right)$, we can invert the F.O.C. condition in (17) to calculate the marginal cost at the fulfillment centers for both logistical modes. Then we can calculate the supply-side moment condition $\mathbf{g}^S(\Gamma)$ as in (22) and (23). We then choose the Γ to minimize the supply-side moment condition $\mathbf{g}^S(\Gamma)' W \mathbf{g}^S(\Gamma)$.

We estimate the standard error of the parameters following Berry et al. (1995). Specifically, the covariance matrix of the parameters is

$$(\Pi'\Pi)^{-1}\Pi' V \Pi (\Pi'\Pi)^{-1}, \tag{25}$$

where $\Pi \equiv \frac{\partial E(G)}{\partial x'}$ is the moment derivatives w.r.t. parameters; and $V = E(Z'u'uZ)$.

6.4 Estimation Results

Table 8 presents our baseline estimation results. First, the store substitution heterogeneity confirms our findings in Section 3.2. We find online stores are more substitutable with one another ($\theta_{on} > 1$), as are stores belonging to the same retailer firm ($\theta_f > 1$). We also find demographic preferences are heterogeneous: poor households are more price-sensitive and distance-averse,

and they have lower perceptions of online store quality. We find that the majority of the online expenditure difference between the rich and the poor comes from households' perceptions of online store quality rather than price and distance concerns.

Table 8: Estimation Results

Parameters	Meanings	Values	s.e
Demand (substitution heterogeneity)			
θ_{on}	Online Nesting Elast.	1.69	0.024
θ_f	Within-Firm Nesting Elast.	1.98	0.05
λ	Cross-nest Weight	0.74	0.022
Demand (demographic heterogeneity)			
α_{rich}	High Income Price Elast.	-0.37	0.19
α_{poor}	Low Income Price Elast.	-0.41	0.20
β_{rich}	High Income Distance Elast.	-0.23	0.003
β_{poor}	Low Income Distance Elast.	-0.26	0.003
Supply (logistics costs)			
γ_{online}	Online Logistic Elast.	0.204	0.202
$\gamma_{offline}$	Offline Logistic Elast.	0.147	0.009

Note: The estimation results include many store-level qualities for each demographic group.

Given our estimates, we can now use dominance analysis to gauge the relative importance of independent variables (i.e., prices, taxes, distance, and quality) in an estimation model based on their contributions to the overall model fit. We calculate the fitted value of the mean utility as

$$\widehat{\delta^d_{j,z}} = \widehat{Q_{j,d}} + \widehat{\alpha_d \ln p_j} + \alpha_d \widehat{\ln(1+t_{j,z})} + \widehat{\alpha_d \ln d_{j,z}}. \tag{26}$$

We then run a dominance analysis to rank the factors by their contributions to total variations in the mean utility. Table 9 presents the results. The results suggest that store quality is the most important factor determining consumer utility, followed by distance, price, and tax considerations.

Our supply-side estimates suggest that spatial frictions matter for logistics even though online shopping can eliminate spatial frictions from the demand side. We find logistics costs to fulfill online orders are 40% higher than those to fulfill offline orders, conditional on the same distance. Therefore, online prices are higher, reflecting their higher costs.

We estimate the online-online substitution elasticity to be -1.8, which is aligned with some existing estimates. Einav et al. (2014) treats tax rates as surprises to consumers when checking

Table 9: Dominance Analysis of Factors that Affect Consumer Purchasing Decisions

$\ln\delta$	Dominance Stat.	Standardized Domin. Stat.	Ranking
Store Quality	0.0356	0.7208	1
Distance	0.0086	0.1733	2
Price	0.0050	0.1014	3
Tax	0.0002	0.0045	4

out and estimates online substitution elasticities around -2. Houde et al. (2021) use Amazon's entry prior to the Wayfair Decision and estimate an average substitution elasticity of -1.5 among all choices. However, all these estimates are significantly smaller than the calibrated elasticities of -8 in Dolfen et al. (2019) and -12 in Couture (2013). These larger elasticities are derived by converting the distance elasticities with gas prices. Namely, $\alpha = \beta / P_{\$/m}$. However, if the disutility from distance includes time value beyond gas price, $p_{\$/m}$ will be underestimated and lead to an overstatement of α.

6.5 Robustness Checks

We are conducting the following additional robustness checks and will update our results in the next version of this paper.

Zero Market Shares

One concern in our handling of distance elasticity is that if a store has no receipts in our data, then it does not belong to a consumer choice set. This could lead to a positive selection of $\tilde{\xi}_{j,z}$ into our estimation sample and an underestimation of the distance elasticity β. We address this concern by including zero market shares and estimating within-firm offline market shares using the PPML estimator proposed by Silva and Tenreyro (2006). According to our demand

system, the market share of offline store j belonging to firm f is[22]

$$
\begin{aligned}
Sales_{j,z} &= \frac{\left(\delta_{j,z}\right)^{\theta_f} \Phi_f^{1/\theta_f - 1}}{\Phi_z} Y_z \\
&= \exp(\underbrace{\theta_f \left(A_j - \alpha \ln p_j (1+\tau_j)\right)}_{\ln F_j} + \underbrace{(1/\theta_f - 1)\ln \Phi_f + \ln Y_z - \ln \Phi_z}_{\ln F_{f,z}} - \underbrace{\theta_f \beta}_{\gamma} \ln d_{j,z} + \theta_f \xi_{j,z}) \\
&= F_j F_{f,z} \exp(-\gamma \ln d_{j,z} + \epsilon_{j,z}). \qquad (27)
\end{aligned}
$$

Then we can treat $Sales_{j,z}$ with zeros as count data and consistently estimate $\gamma = \theta_f \beta$ in (27) using the PPML estimator with two-way fixed effects. Table A.4 presents our estimation results.

Other BLP Instruments

Another concern related to one of our instruments, $N^{rival}_{on,j,z,t}$, which is the number of rival online stores in the market, is the possibility of measurement error because of finite sampling. We address this concern by introducing another set of BLP instruments that are not subject to this measurement concern. This set is the sum of non-rival and rival online stores' tax rates. $T^{rival}_{on,j,z,t} \equiv \sum_{j'} \tau_{j',z,t}$, $T^{non\text{-}rival}_{on,f,z,t} \equiv \sum_{j''} \tau_{j'',z,t}$ where $j'(j'')$ are rival (non-rival) online stores of store j, and $\tau_{\cdot,z,t}$ is the sales tax rate for ZIP z in month t.

Offline-channel Nest

One may be interested in whether stores belonging to the offline channel are more substitutable. We do not impose such a strong prior on an offline-channel nest because we cannot learn this from online tax shocks. However, our demand system can incorporate that research interest and identify relevant elasticities by leveraging the spatial variations of the offline stores. To introduce the offline-channel nest, we introduce two additional parameters: $\theta_{off}, \lambda_{off}$. θ_{off} describes how substitutable offline-channel nests are, and λ_{off} is a splitting parameter for all offline stores governing the weights between an offline nest and a firm nest. The generalized market shares thus are

$$
S_{j,z} = \frac{\left(\lambda_c \delta_{j,z}\right)^{\theta_c} \Phi_{c,z}^{1/\theta_c - 1} + \left((1-\lambda_c)\,\delta_{j,z}\right)^{\theta_f} \Phi_{f,z}^{1/\theta_f - 1}}{\Phi_z}, \qquad (28)
$$

[22] We suppress time and demographic subscripts t,d to facilitate discussion.

where $c \in \{on, off\}$ is the channel of store j. To identify $\theta_{off}, \lambda_{off}$, we need two additional BLP instruments: $N^{rival}_{off,j,z,t}, N^{non\text{-}rival}_{off,j,z,t}$. These are the number of rival and non-rival offline stores in the market, respectively. We estimate $\theta_{offline} = 1.60, \lambda_{offline} = 0.062$, which suggests the offline stores are mostly close substitutable with other offline stores within the same retail chain. Therefore, our baseline model is a good approximation of the economy.

COVID Shocks and Online Shopping

One may wonder if COVID shocks bias our estimation results since our data sample includes data from 2020 and 2021. Our specification can accommodate this concern because we estimate the model with demographic-store-time fixed effects, capturing all store-level quality changes due to COVID. Our identification assumption is that consumers' preferences for prices and spatial frictions are stable over time. We can test this assumption by estimating (27) for different time periods. Table A.5 presents our estimation results.

7. Counterfactuals

The demand and supply model is a laboratory to evaluate many counterfactual online policies. We perform two exercises in this section. In Section 7.1, we evaluate the Wayfair Decision's effects on consumer welfare. In Section 7.2, we remove all the online stores and decompose the welfare effects into convenience, variety, and pro-competitive effects. We also characterize the distributional effect of e-commerce.

To conduct these counterfactuals, we fix the demand and supply shocks $(\xi_{j,z,t}, \eta_{f,d,t})$ and store-level fixed quality $A_{j,t}$ and change only the policy-relevant parameters. We show that conditional on observed market shares, the counterfactual results depend only on the key elasticities in our model. Therefore, we solve for the counterfactual equilibrium using exact-hat algebra (à la Dekle et al. (2007)). We then express the welfare effects of counterfactual in terms of compensating variation,

$$CV = \frac{\Phi_z}{\Phi_{z'}} - 1, \tag{29}$$

where Φ_z ($\Phi_{z'}$) is the price index in the initial (counterfactual) economy defined as in (13).

7.1 The Welfare Effects of the Wayfair Decision

As a first exercise, we calculate the welfare effects of the Wayfair Decision. To do so, we start our economy in 2021 and hold all parameters constant with the exception of the online tax rate, which we set to the pre-Wayfair-era rate. According to our findings in Section 3, retailers' pricing does not respond to the Wayfair Decision's tax shocks. Therefore, in this exercise, we fix the prices as of 2021. In this case, only demand-side parameters $(\alpha_{rich}, \alpha_{poor}, \theta_{on}, \theta_f, \lambda)$ matter. We calculate the compensating variation of consumer welfare due to this tax policy change. Figure 12 shows the results.

Figure 12: Welfare Effects of the Wayfair Decision

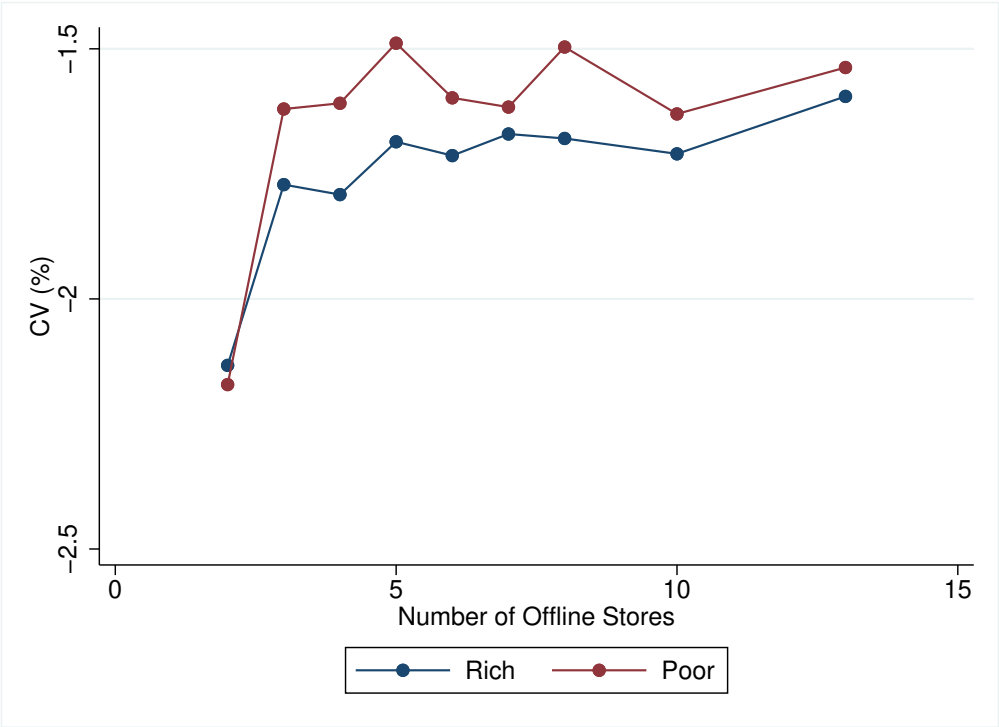

Note: We calculate the compensating variation of consumer welfare when changing the sales tax rates to their pre-Wayfair rates but holding prices fixed. Then we binned-scatter plot the average welfare change against urban density.

We find that in the counterfactual scenario, richer households living in rural areas benefit the most since they are willing to sacrifice more of their shopping budget to prevent the counterfactual from happening. Therefore, we conclude that the Wayfair Decision has more negative effects on households with higher incomes and households in rural areas than it does on other households.

7.2 Welfare Gains from Online Shopping

In this section, we evaluate the overall and distributional effects of rising e-commerce on consumer welfare. We start our economy in 2021 and roll out our counterfactual in three steps. First, we assume online shopping causes consumers to incur travel costs equal to the average offline shopping distance. This allows us to measure the gains from convenience. Then, we remove all the online stores but hold everything else, including prices, fixed. We refer to the welfare change in this step as gains from online variety. Third, we allow the remaining firms to optimally adjust their prices according to our supply model. The welfare change in this step is the gains from the pro-competitive effect.

We use exact-hat algebra to solve for the counterfactual price equilibrium. We define the price change of each firm-channel-region as $\hat{p} \equiv \frac{p_1}{p_0}$. We first guess a vector of prices changes $\hat{\mathbf{p}}$ and solve for the new market share $\mathbf{S}'$ and Jacobian Δ' using (12). Then we update the prices using the F.O.C. condition until these two systems converge:

$$\frac{p^{new}_{f,c,r}}{p_{f,c,r}} = -\frac{\sum_s \sum_z w'_{s,z} d^{\gamma_c}_{c,s,z}(1-\Delta^{s,c'}_{f,c})}{\sum_s \sum_z w'_{s,z} \Delta^{s,z'}_{f,c}}. \tag{30}$$

We find that the average gain from convenience is 5% and the average gain from variety is 9%, both in terms of consumers' shopping budget. When we allow all remaining firms to adjust their prices optimally, we find that offline sellers increase their prices by around 13%. We further calculate that the pro-competitive effect is 3% of consumers' shopping budget. Therefore, we conclude that for the average consumer, the welfare gains from rising e-commerce are 17%.

Dolfen et al. (2019) estimates the welfare gains from e-commerce to be around 1%, which is significantly lower than our results. We reconcile the difference by noting the following: (1) the pet-food online share (40%) is significantly larger than other sectors' average (8 %); and (2) unlike Dolfen et al. (2019)'s estimates, our estimates suggest the demand for online stores is inelastic in price. Both of these factors contribute to our large online welfare gains.

To gauge the distributional effect of rising e-commerce, we bin all ZIP codes into 20 groups and calculate the mean welfare change of rich and poor consumers within each bin. We present the results in 13. We find rich households and households in rural areas suffer more than others when online stores are removed. Therefore, we conclude that the rise of e-commerce has reduced consumption inequality between rural and urban areas but increased consumption

inequality between the rich and the poor.

Figure 13: Distributional Effect of E-commerce

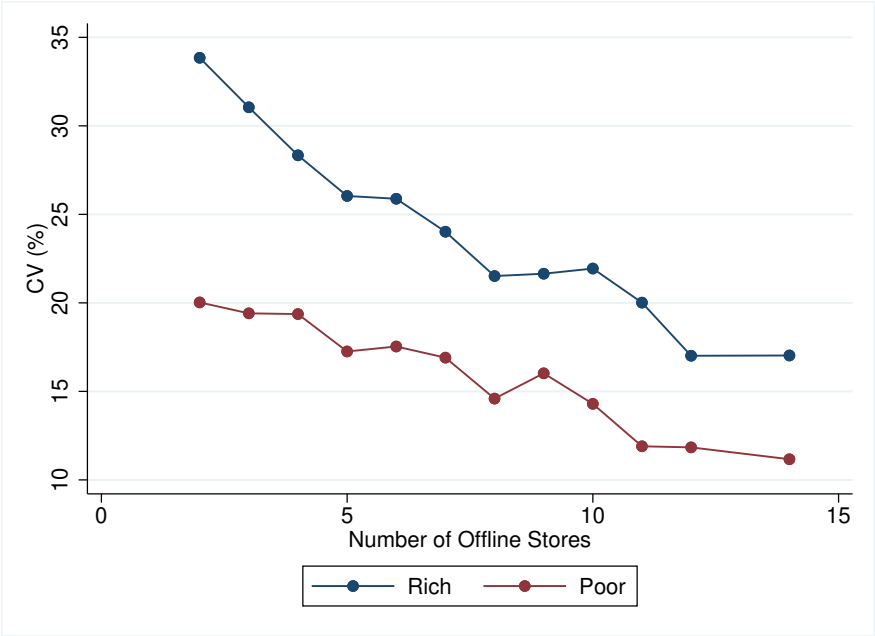

Note: We calculate the compensating variation of consumer welfare for rich and poor households after removing all online stores for each ZIP code. Then we binned-scatter plot the average welfare change against urban density.

8. Conclusion

In this paper, we studied how the rise of e-commerce has reshaped overall consumer welfare and consumption inequality in the presence of retail oligopoly. We used shopping receipts data to document new stylized facts about online retail markets. We found that households in rural areas with higher incomes tend to shop online more than other households. Then we leveraged an exogenous tax shock caused by the Supreme Court's Wayfair Decision to learn about online demand and firm pricing response. Our results indicate that the substitution among stores is not homogeneous: stores belonging to the same channel or the same retailer firm are more substitutable.

Motivated by those facts, we developed a pet-food demand and supply model to estimate consumer shopping preferences and retailer logistics costs. The estimation results suggest that although online stores greatly reduce the spatial friction for the demand side, they do so at a

higher logistics cost than incurred in the offline shipping mode, which in turn raises online prices. We then used the model to evaluate welfare gains and measure the distributional effects of rising e-commerce on different consumers. For the average consumer, the gains from price competition are quantitatively as important as the gains from convenience. Our results suggest that the rise of e-commerce has reduced consumption inequality between rural and urban areas but increased consumption inequality between the rich and the poor.

References

Adams, Brian and Kevin R. Williams, "Zone Pricing in Retail Oligopoly," *American Economic Journal: Microeconomics*, February 2019, *11* (1), 124–156.

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy*," *The Quarterly Journal of Economics*, August 2014, *129* (3), 1085–1140.

Anderson, James E and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, February 2003, *93* (1), 170–192.

Arcidiacono, Peter, Paul B. Ellickson, Carl F. Mela, and John D. Singleton, "The Competitive Effects of Entry: Evidence from Supercenter Expansion," *American Economic Journal: Applied Economics*, July 2020, *12* (3), 175–206.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *63* (4), 841.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficient Estimation," April 2022. arXiv:2108.12419 [econ].

Butters, R. Andrew, Daniel W. Sacks, and Boyoung Seo, "How Do National Firms Respond to Local Cost Shocks?," *American Economic Review*, May 2022, *112* (5), 1737–1772.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, December 2021, *225* (2), 200–230.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, August 2009, *99* (4), 1145–1177.

Couture, Victor, "Valuing the Consumption Benefits of Urban Density," 2013, p. 70.

_ , Benjamin Faber, Yizhen Gu, and Lizhi Liu, "Connecting the Countryside via E-Commerce: Evidence from China," *American Economic Review: Insights*, March 2021, *3* (1), 35–50.

Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales, "How Segregated Is Urban Consumption?," *journal of political economy*, 2019, p. 55.

de Chaisemartin, Clément and Xavier D'Haultfoeuille, "Difference-in-Differences Estimators of Intertemporal Treatment Effects," March 2022, p. w29873.

Dearing, Adam, "Estimating structural demand and supply models using tax rates as instruments," *Journal of Public Economics*, January 2022.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced Trade," *American Economic Review*, April 2007, *97* (2), 351–355.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform Pricing in U.S. Retail Chains*," *The Quarterly Journal of Economics*, November 2019, *134* (4), 2011–2084.

Diamond, Rebecca and Enrico Moretti, "Where is Standard of Living the Highest? Local Prices and the Geography of Consumption," Technical Report w29533, National Bureau of Economic Research, Cambridge, MA December 2021.

Dolfen, Paul, Liran Einav, Peter Klenow, Benjamin Klopack, Jonathan Levin, Laurence Levin, and Wayne Best, "Assessing the Gains from E-Commerce," Technical Report w25610, National Bureau of Economic Research, Cambridge, MA February 2019.

Donaldson, Dave and Richard Hornbeck, "Railroads and American Economic Growth: A "Market Access" Approach*," *The Quarterly Journal of Economics*, May 2016, *131* (2), 799–858.

Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan, "Sales Taxes and Internet Commerce," *American Economic Review*, January 2014, *104* (1), 1–26.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The Return to Protectionism*," *The Quarterly Journal of Economics*, February 2020, *135* (1), 1–55.

Fally, Thibault, "Structural gravity and fixed effects," *Journal of International Economics*, September 2015, *97* (1), 76–85.

Fan, Jingting, Lixin Tang, Weiming Zhu, and Ben Zou, "The Alibaba effect: Spatial consumption inequality and the welfare gains from e-commerce," *Journal of International Economics*, September 2018, *114*, 203–220.

Feenstra, Robert C., "Exact Hedonic Price Indexes," *The Review of Economics and Statistics*, November 1995, *77* (4), 634.

Feng, Guiyun, Xiaobo Li, and Zizhuo Wang, "On substitutability and complementarity in discrete choice models," *Operations Research Letters*, January 2018, *46* (1), 141–146.

Gentzkow, Matthew, "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, May 2007, *97* (3), 713–744.

Goolsbee, Austan, "In a World Without Borders: The Impact of Taxes on Internet Commerce*," *The Quarterly Journal of Economics*, May 2000, *115* (2), 561–576.

Handbury, J. and D. E. Weinstein, "Goods Prices and Availability in Cities," *The Review of Economic Studies*, January 2015, *82* (1), 258–296.

Handbury, Jessie, "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," *Econometrica*, 2021, *89* (6), 2679–2715. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA11738.

Hausman, Leonard, and Zona, "Competitive Analysis with Differenciated Products," *Annales d'Économie et de Statistique*, 1994, (34), 159.

Hossain, Mallick, "A world without borders revisited: Impact of online sales tax collection on shopping and search," *Journal of Economics & Management Strategy*, February 2022, *31* (1), 48–63.

Houde, Jean-François, Peter Newberry, and Katja Seim, "Economies of Density in E-Commerce: A Study of Amazon's Fulfillment Center Network," April 2021.

Huang, Yufeng and Bart J. Bronnenberg, "Consumer Transportation Costs and the Value of E-commerce: Evidence from the Dutch Apparel Industry," SSRN Scholarly Paper ID 3596460, Social Science Research Network, Rochester, NY December 2021.

Jia, Panle, "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 2008, *76* (6), 1263–1316.

Jo, Yoon J., Misaki Matsumura, and David E. Weinstein, "The Impact of E-Commerce on Relative Prices and Consumer Welfare," November 2019.

Kroft, Kory, Jean-William Laliberté, René Leal-Vizcaíno, and Matthew Notowidigdo, "Salience and Taxation with Imperfect Competition," Technical Report w27409, National Bureau of Economic Research, Cambridge, MA June 2020.

McFadden, Daniel, "Econometric Models for Probabilistic Choice Among Products," *The Journal of Business*, 1980, *53* (3), S13–S29. Publisher: University of Chicago Press.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen Redding, "The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data," Technical Report w28497, National Bureau of Economic Research, Cambridge, MA February 2021.

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, March 2001, *69* (2), 307–342.

Papola, Andrea, "Some developments on the cross-nested logit model," *Transportation Research Part B: Methodological*, November 2004, *38* (9), 833–851.

Relihan, Lindsay, "Is Online Retail Killing Coffee Shops? Estimating the Winners and Losers of Online Retail Using Customer Transaction Microdata," *SSRN Electronic Journal*, 2022.

Rossi-Hansberg, Esteban, Nicholas Trachter, Ezra Oberfield, and Pierre-Daniel Sarte, "Plants in Space," 2020, p. 68.

Seim, Katja, "An empirical model of firm entry with endogenous product-type choices," *The RAND Journal of Economics*, September 2006, *37* (3), 619–640.

Seo, Jaeeun and Ryungha Oh, "What Causes Agglomeration of Services? Theory and Evidence from Seoul," 2022.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, November 2006, *88* (4), 641–658.

Song, Jintaek, "MPC of SNAP and its Welfare Implication," 2022.

Sullivan, Micheal, “Price controls in a multi-sided market,” 2022.

Zoutman, Floris T., Evelina Gavrilova, and Arnt O. Hopland, “Estimating Both Supply and Demand Elasticities Using Variation in a Single Tax Rate,” *Econometrica*, 2018, *86* (2), 763–771.

Appendix

A. Data Appendix

A.1 Numerator Data: Receipts Example

We show examples of the original receipts panelists upload in this section.

(a) Example of Offline Shopping Receipts

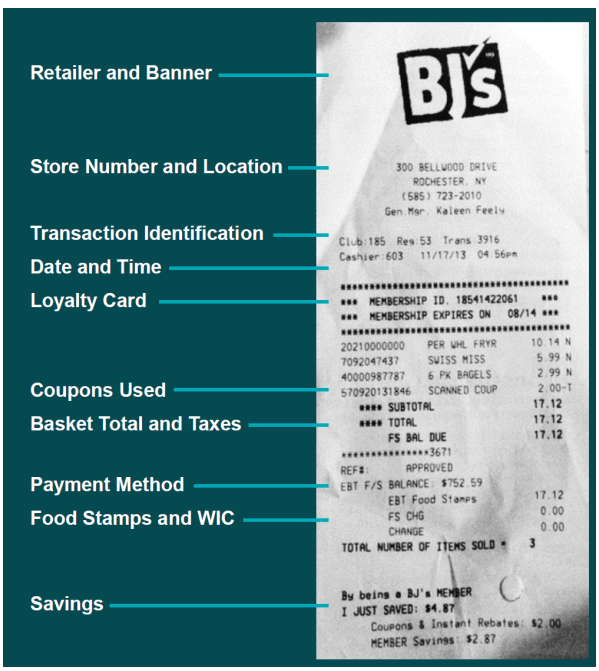

(b) Example of Online Shopping Receipts

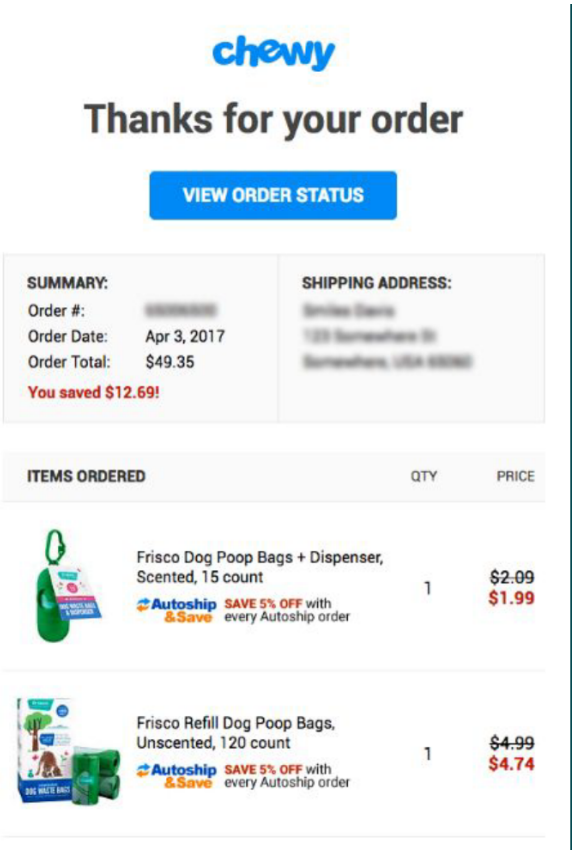

Note: We show examples of original receipts collected from the Numerator Panelist. The left panel is an example of offline receipts, from which we learn the retailer names and store addresses, and information on items bought, including prices, quantities, item descriptions, payment methods, subtotal, and taxes. The right panel is an example of online receipts, from which we learn retailer names, information on items bought including prices, quantities, and item description, subtotal, and taxes.

A.2 Compare Numerator with existing datasets

Compared with online transaction data, such as Comscore and Forrester, Numerator data document consumers' omnichannel behaviors, allowing us to learn their store preferences over both channels. Compared with visa transaction data used in Dolfen et al. (2019) and Relihan (2022), Numerator has detailed prices and quantity information, which we will use to estimate consumer price elasticity. Probably Nielsen consumer panel is the most comparable one. However, Nielsen data has the following shortages: 1) Retailer identities are masked, and we can not link online stores with offline stores for the same firm, which prevents us from learning online-offline cannibalization. 2) Most store addresses are missing and only reported to the zip code level, which will cause huge measurement errors in the estimation of distance elasticity, a salient feature to understand in contrast with online shopping. 3) Panelists' barcode scanning suffers from missing items and reporting bias. Their self-reporting total spending may lead to large measurement errors to calculate taxes, so we can not use it to study the Wayfair Decision. Therefore, Numerator data is a better fit for our purposes.

Table A.1: Comparison of Numerator Data Set to Other Data Sets

	Omnichannel	Price	Tax	Retailer info
Numerator	✓	✓	✓	✓
Nielsen	✓	✓	✗	✗
Comscore	✗	✓	✓	✓
Forrest	✗	✗	✗	✗
Credit Card	✓	✗	✗	✓

A.3 Summary Statistics of All Sectors

Table A.2: Summary Statistics of All Sectors in Chicago Kilts Center Archive of Numerator Data (2017-2021)

Year	2017		2018		2019		2020		2021	
Sector	(m$)	Online	(m$)	Online	(m$)	Online	(m$)	Online	(m$)	Online
Apparel	155.0	64%	173.0	55%	147.0	48%	96.4	46%	97.6	43%
Automotive	28.9	74%	32.7	68%	28.5	60%	20.6	55%	22.1	47%
Baby	60.4	54%	72.0	46%	65.1	40%	50.8	36%	55.5	35%
Books	29.8	87%	28.6	84%	21.5	79%	15.1	71%	11.6	57%
Electronics	216.0	83%	221.0	78%	177.0	73%	140.0	69%	133.0	64%
Entertainment	11.4	53%	10.7	45%	8.2	40%	5.0	37%	4.1	33%
Grocery	861.0	4%	1160.0	4%	1240.0	4%	1390.0	4%	1540.0	8%
Health & Beauty	222.0	38%	280.0	34%	284.0	30%	266.0	28%	279.0	27%
Home & Garden	182.0	69%	215.0	63%	191.0	57%	169.0	55%	187.0	50%
Household	110.0	21%	147.0	18%	158.0	16%	167.0	15%	167.0	17%
Office	30.4	61%	35.9	57%	31.4	52%	25.9	51%	26.8	46%
Party & Occasions	62.8	28%	75.6	28%	73.6	24%	61.6	24%	66.2	23%
Pet	71.7	39%	95.8	36%	100.0	34%	95.8	31%	121.0	40%
Quick Serve Restaurant	18.5	96%	31.6	98%	17.6	96%	7.0	91%	5.7	79%
Restaurant	0.1	2%	0.1	3%	0.1	4%	0.1	6%	0.2	6%
Sports	35.8	82%	36.7	75%	28.8	68%	23.3	68%	21.3	54%
Tobacco Products	17.0	1%	21.6	1%	24.4	1%	29.2	1%	37.6	1%
Tools & Improvement	40.5	76%	47.6	72%	42.1	68%	39.0	66%	40.6	60%
Toys	76.8	64%	84.4	58%	69.5	52%	57.7	51%	59.7	44%

Note: Source: Numerator Data. Notice that since Chicago Kilts Archive miss significantly offline stores with special durable goods (e.g. Home Depot, Best Buy, IKEA, Apple), the online shares could be significantly overstated.

A.4 Representativeness of Numerator Data

In this section, we compare Numerator Data with US Census in Table A.3.

Table A.3: Comparison of Numerator Panelist Demographics and US Census

	Numerator (%)	Census (%)
Age		
18-24	2.4	5.1
25-34	20.5	16.1
35-44	28.2	17.0
45-54	23.9	18.9
55-64	16.4	18.9
Over 65	8.7	24
Income		
Under $20k	13.5	18
$20k-40k	20.4	20.3
$40k-60k	19.1	15.9
$60k-80k	14.4	12.4
$80k-125k	20.7	16.9
Over $125k	12.0	16.5
Ethnicity		
White/Caucasian	69.5	67.6
Black or African American	8.5	13.2
Hispanic/Latino	10.5	13
Asian	8.1	5
Other	3.5	1.2
Census Division		
Northeast	18.6	18.0
Midwest	22.9	22.4
South	39.4	37.4
West	19.0	22.2
Urbanicity		
Rural	27.1	22.2
Suburban	40.5	39.5
Urban	32.4	38.3

Note: Source: 2019 Numerator and Census Data.

A.5 Fulfillment Centers of Major Pet Food Retailers

We visualized the distribution of fulfillment centers of the major retailers in the following graph.

Figure A.2: Fulfillment centers of major pet food retailers

Note: The distribution centers of the major pet food retailers in 2021. Source: Infogroup data.

A.6 More Narrative Evidence on the Effect of the Wayfair Decision

In this section, we provide two additional pieces of narrative evidence on the effect of the Wayfair Decision on Chewy.com.

1. "Changes in the tax treatment of companies engaged in e-commerce may **adversely affect** the commercial use of our website and mobile applications and our financial results." — CHEWY, INC. FORM 10-K

2. "I just ended my long relationship with Chewy.com. Not only ..., but now Chewy **adds another $13.71 in tax, which they never charged before**. It actually **costs more** to buy from Chewy than it costs to **run 1 mile down the street and buy it at the tiny pet shop.**" — User comments from Chewy.com Facebook pages

A.7 Marketplace Sellers Imputation

In the following graph, we give an example of what is exactly a marketplace seller on Amazon.com is subject to the Wayfair Decision.

The Wayfair Decision will only affect the Amazon/Walmart marketplace sellers rather than Amazon.com and Walmart.com. One challenge we face is we can not directly distinguish marketplace sellers in our receipts data. Therefore, we impute item marketplace status at the brand-state level. For example, if we observe, before Wayfair Date, most of the transactions containing brand X are tax-free, then we classify brand X as a marketplace brand. We show an example of our imputation process in Figure A.4. Consider two brands – Hills Science Diet and Rocco & Roxie selling on Amazon in California. We classify Rocco & Roxie as marketplace sellers because, before 2019 Oct, most shopping receipts containing at least one item from Rocco & Roxie are significantly below than California sales tax rate of 7.25%. We further validate this classification by the listing status on the Amazon Website. See Appendix A.3.

We acknowledge that our algorithm has the tendency to underestimate the marketplace sellers' market share: if a brand is sold by both Amazon.com and third-party sellers on the Amazon marketplace, our algorithm tends to classify that brand as Amazon.com.

A.8 Jump-detection Algorithm for Sales Tax Shock

As we stated in the previous section, even though each state set up its own deadline for adjustment, each retailer follows its own schedule. Plenty of retailers adjusted their tax scheme ahead of the state deadline. Therefore, the timing of the treatment should at the retailer-state-time level.

We collect Amazon and Walmart marketplace adjustment dates from their websites[23]. Then we develop a jump-detection algorithm to impute the adoption date of Chewy.com using sales tax information from its receipts.

To implement our jump-detection algorithm, we first calculated several key statistics of the effective tax-rate distribution of each retailer's receipts in each state month, including mean, and three quantiles. That will generate a multidimensional time series for each retailer-state-month, which we later input into the python ruptures library to detect any structural breakpoints. We use rupture as the preferred algorithm not only for its fast implementation but also

[23]Amazon marketplace: https://www.amazon.com/gp/help/customer/display.html?nodeId=202211260. Walmart marketplace: https://sellerhelp.walmart.com/seller/s/guide?article=000006444&language=en_US

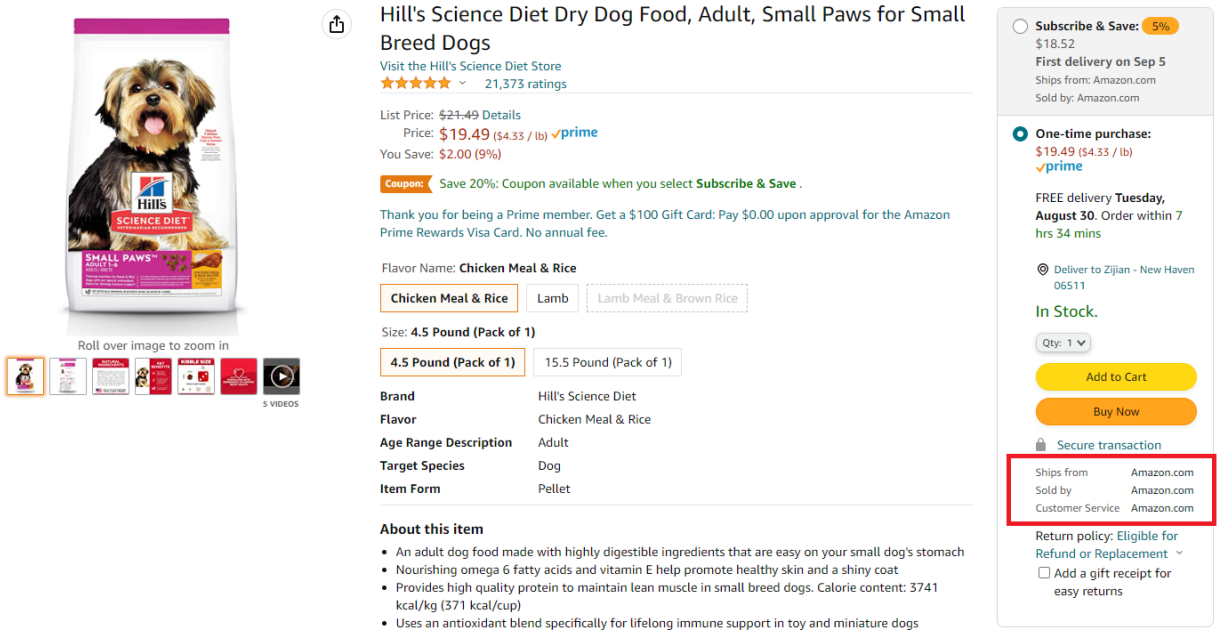

(a) Pet food sold by Amazon.com, apply sales tax all the time, thus not affected by Wayfair Decision.

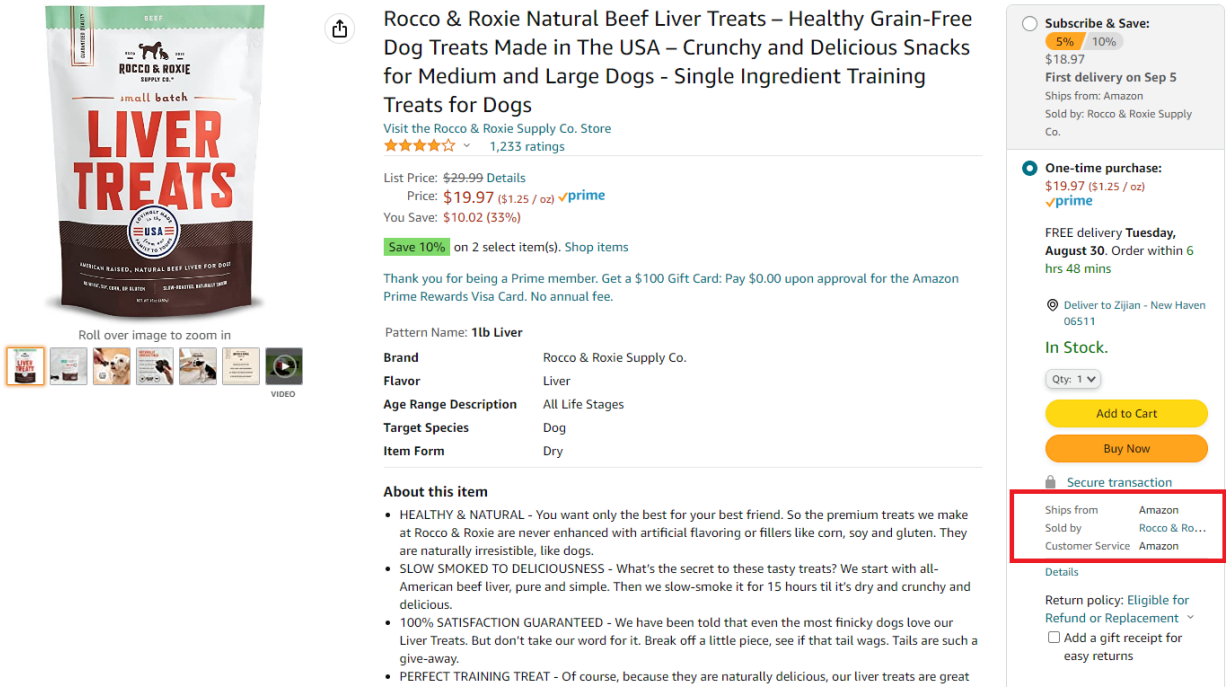

(b) Pet food sold by Amazon Marketplace sellers, used to be tax-free but need to collect sales tax after Wayfair Decisions.

Figure A.3: Pet food sold by the first and third party on Amazon: an example.

Figure A.4: Marketplace seller classification example

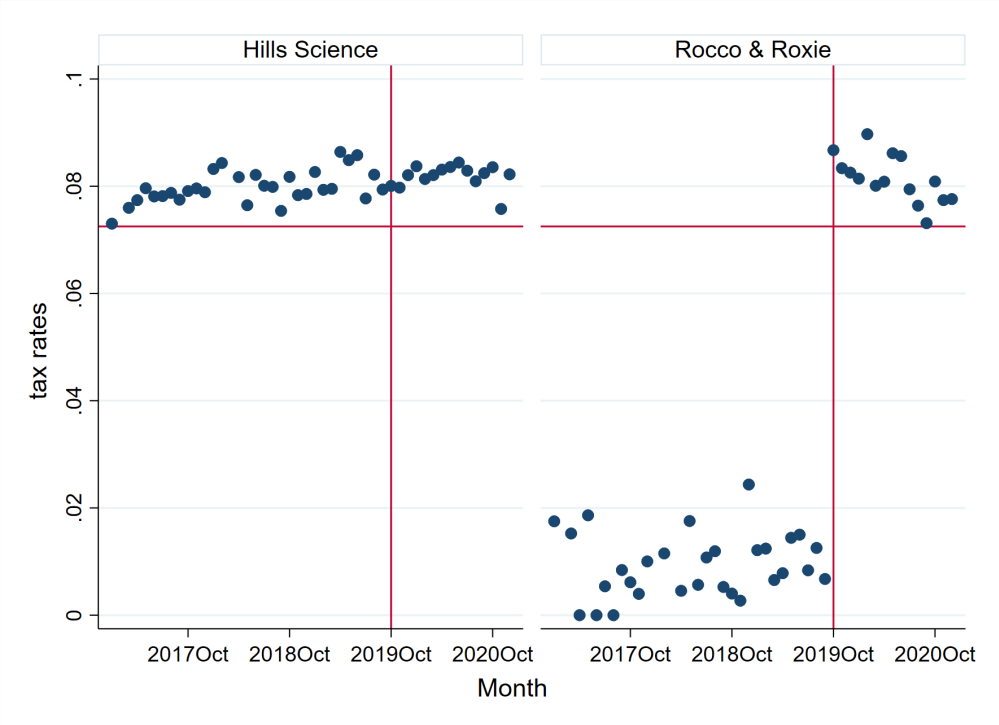

Note: Rocco & Roxie is classified as an Amazon marketplace brand in California since most of the receipts have significantly lower sales tax rates before the Wayfair Date.

due to its capability to handle multiple dimensional time series data. However, ruptures will always output one candidate even if there is no real change. Therefore, we then run a chow-test for each candidate to assure the point is indeed a jump point. Figure (A.5) shows the graphical illustration of what our algorithm is doing. To validate our jump-detection algorithm, we compare the points of time identified by our algorithm with the true date announced by Amazon, as is shown in Figure (A.6).

A.8.1 Store Price Indices

Although online sellers could price-discriminate users based on their IPs, we do not find that's the case in our data. For the pet food market, we find few spatial price variations of the same online retailer across states. We use Walmart's online and offline pricing as an example. For each month, we calculate the Walmart offline stores and walmart.com prices index using (2) in each state and plot their interquartile range (p75 - p25) in Figure . The almost-uniform pricing of online stores raises a challenge for online demand estimation. Since we may not exploit

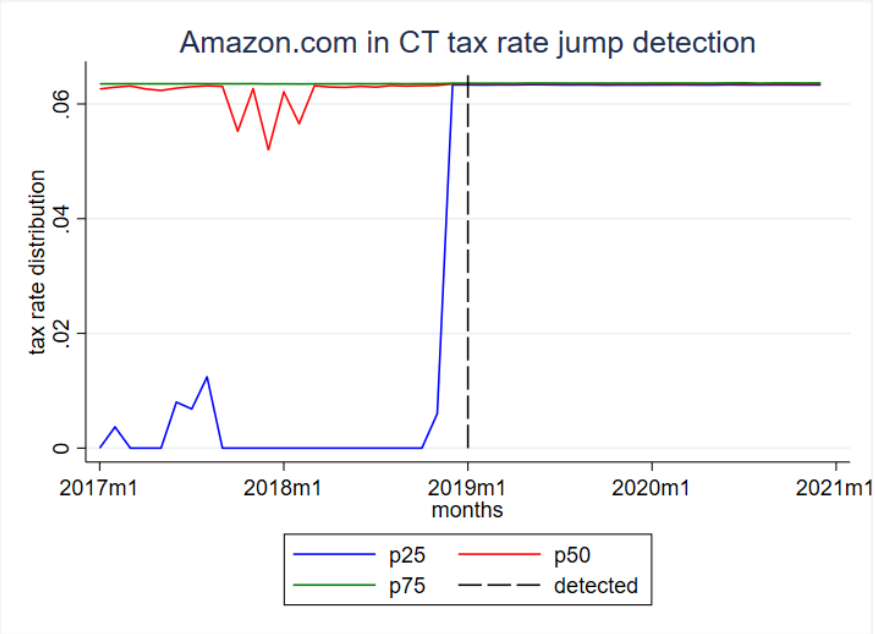

Figure A.5: Online Tax Adoption Time Detection Algorithm

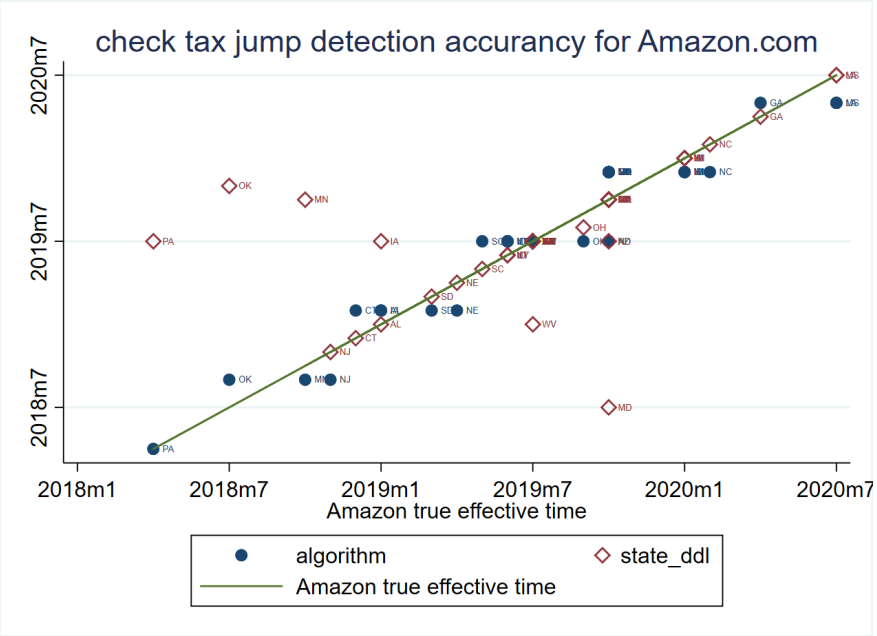

Figure A.6: Accuracy of alogrithm

the across-market price variations as instruments for identification (Hausemen Instruments), we formally discuss this issue in Section 6.2. Therefore, our tax-shock identification strategy becomes particularly helpful for us in learning online demand.

Figure A.7: Walmart Store Prices Distribution (p25-p75)

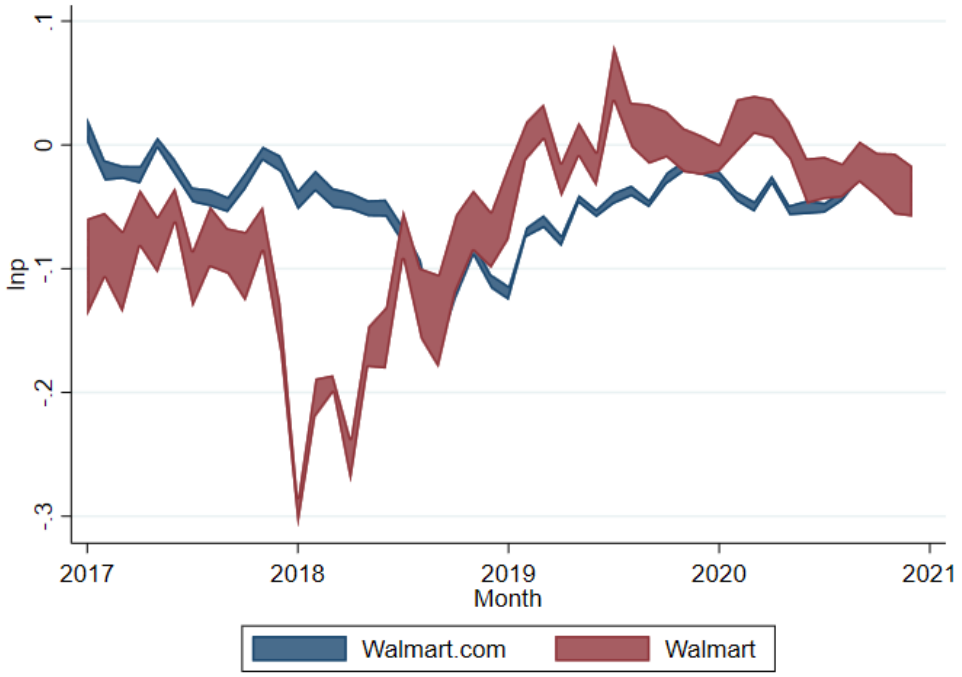

Note: Walmart prices variations: online prices are different from in-store prices, and they have fewer spatial variations.

A.8.2 Evidence on Imperfect Substitutes of retail Chains

We found evidence that store chains are imperfect substitutes. Consumers in the same ZIP typically visit four different Walmart stores. Given our sample has 19k ZIP and only 4461 Walmart Stores, the model assumes perfect substitutions where consumers always choose the closest store is hard to rationalize the data and understate the value of multiple branches to consumer welfare. Therefore, we treat branches belonging to the same retailer as imperfect substitutes in our model.

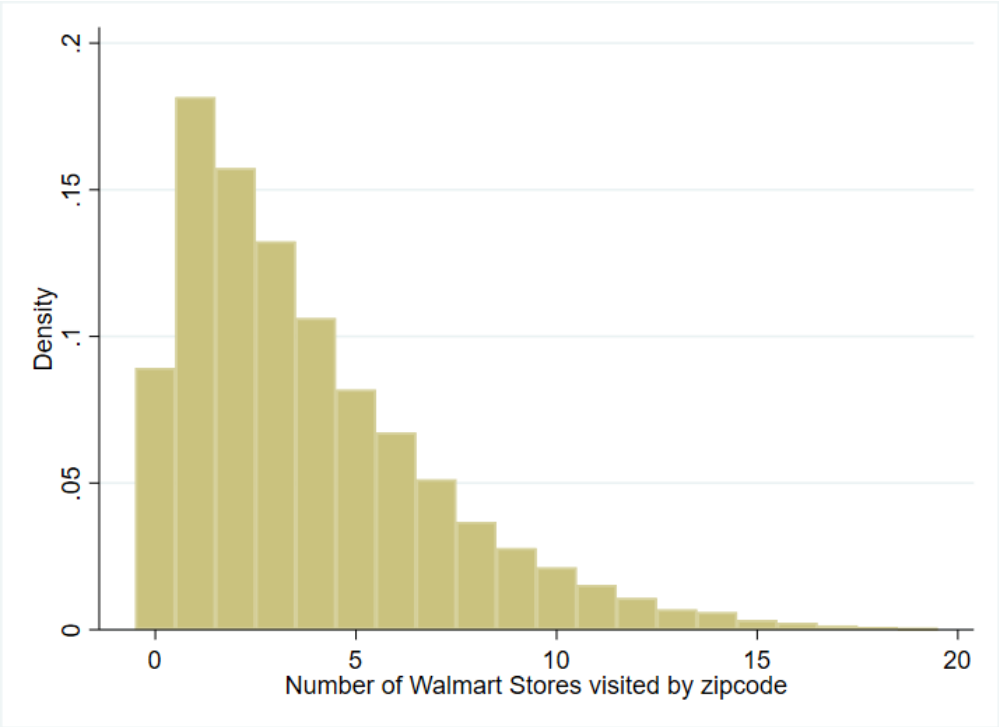

Figure A.8: Distribution of the number of Walmart stores for consumers living in the same ZIP visits.

A.8.3 Event Study

Here we run the did-imputation regression proposed by Borusyak et al. (2022) to address the concern that the simple difference-in-difference event study may impose the wrong weight if there are heterogeneous treatment effects.

A.8.4 Wayfair Effect on all sectors

Figure A.9: did-imputation

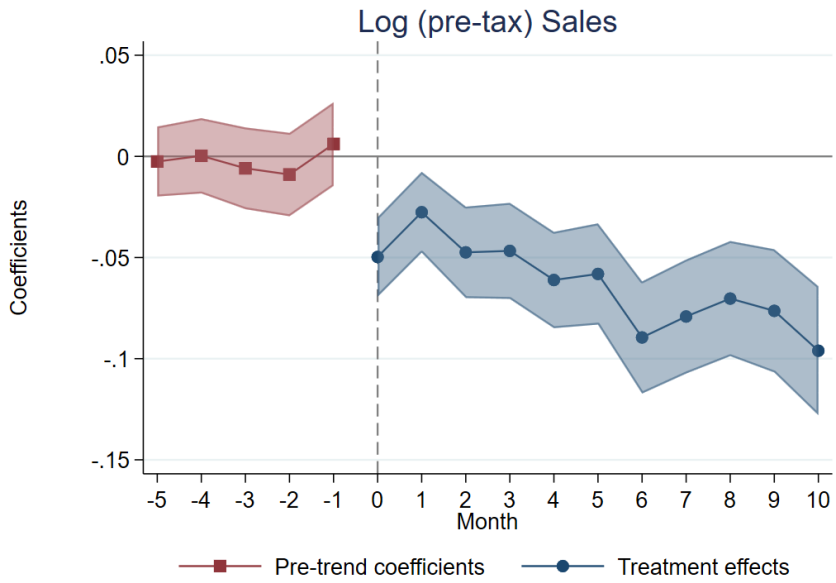

Note: Coefficient plot of the did imputation methods proposed by Borusyak et al. (2022) to address heterogeneous treatment concern

Figure A.10: Placebo Test on Taxable Sectors

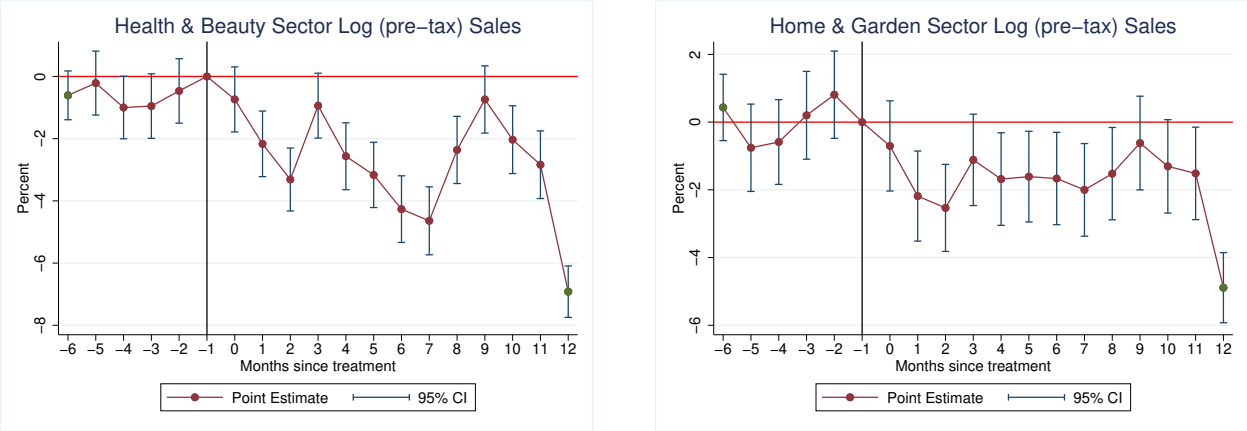

Note: Coefficients plots on β_k of our event study regression on the log of pre-tax sales as in (3). We do find the Wayfair Decision has a significant impact on pre-tax sales in taxable sectors such as home & garden, health & beauty sectors.

Figure A.11: Placebo Test on Non-taxable Sectors

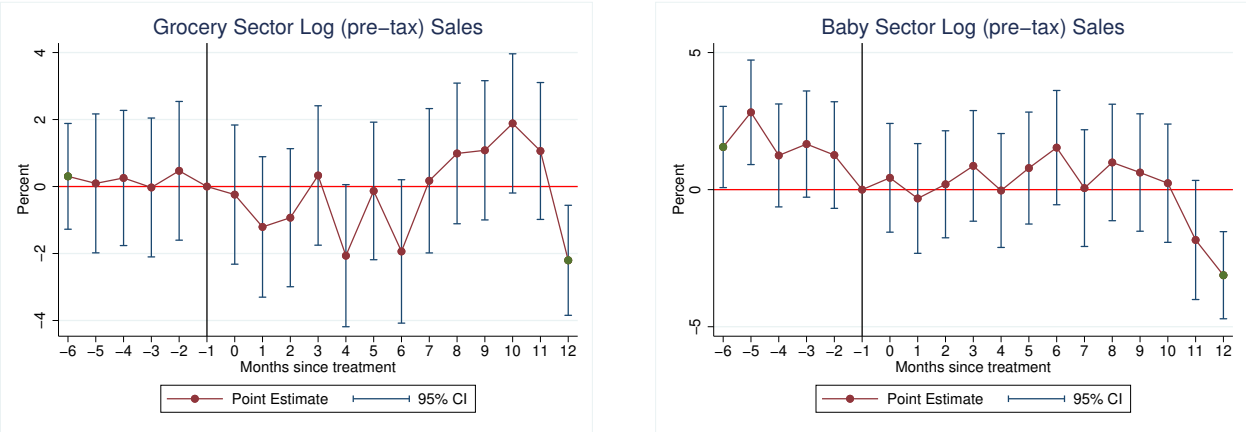

Note: Coefficients plots on β_k of our event study regression on the log of pre-tax sales as in (3). We do not find the Wayfair Decision has an impact on pre-tax sales in non-taxable sectors such as grocery and baby sectors.

Figure A.12: Wayfair Effects over all sectors

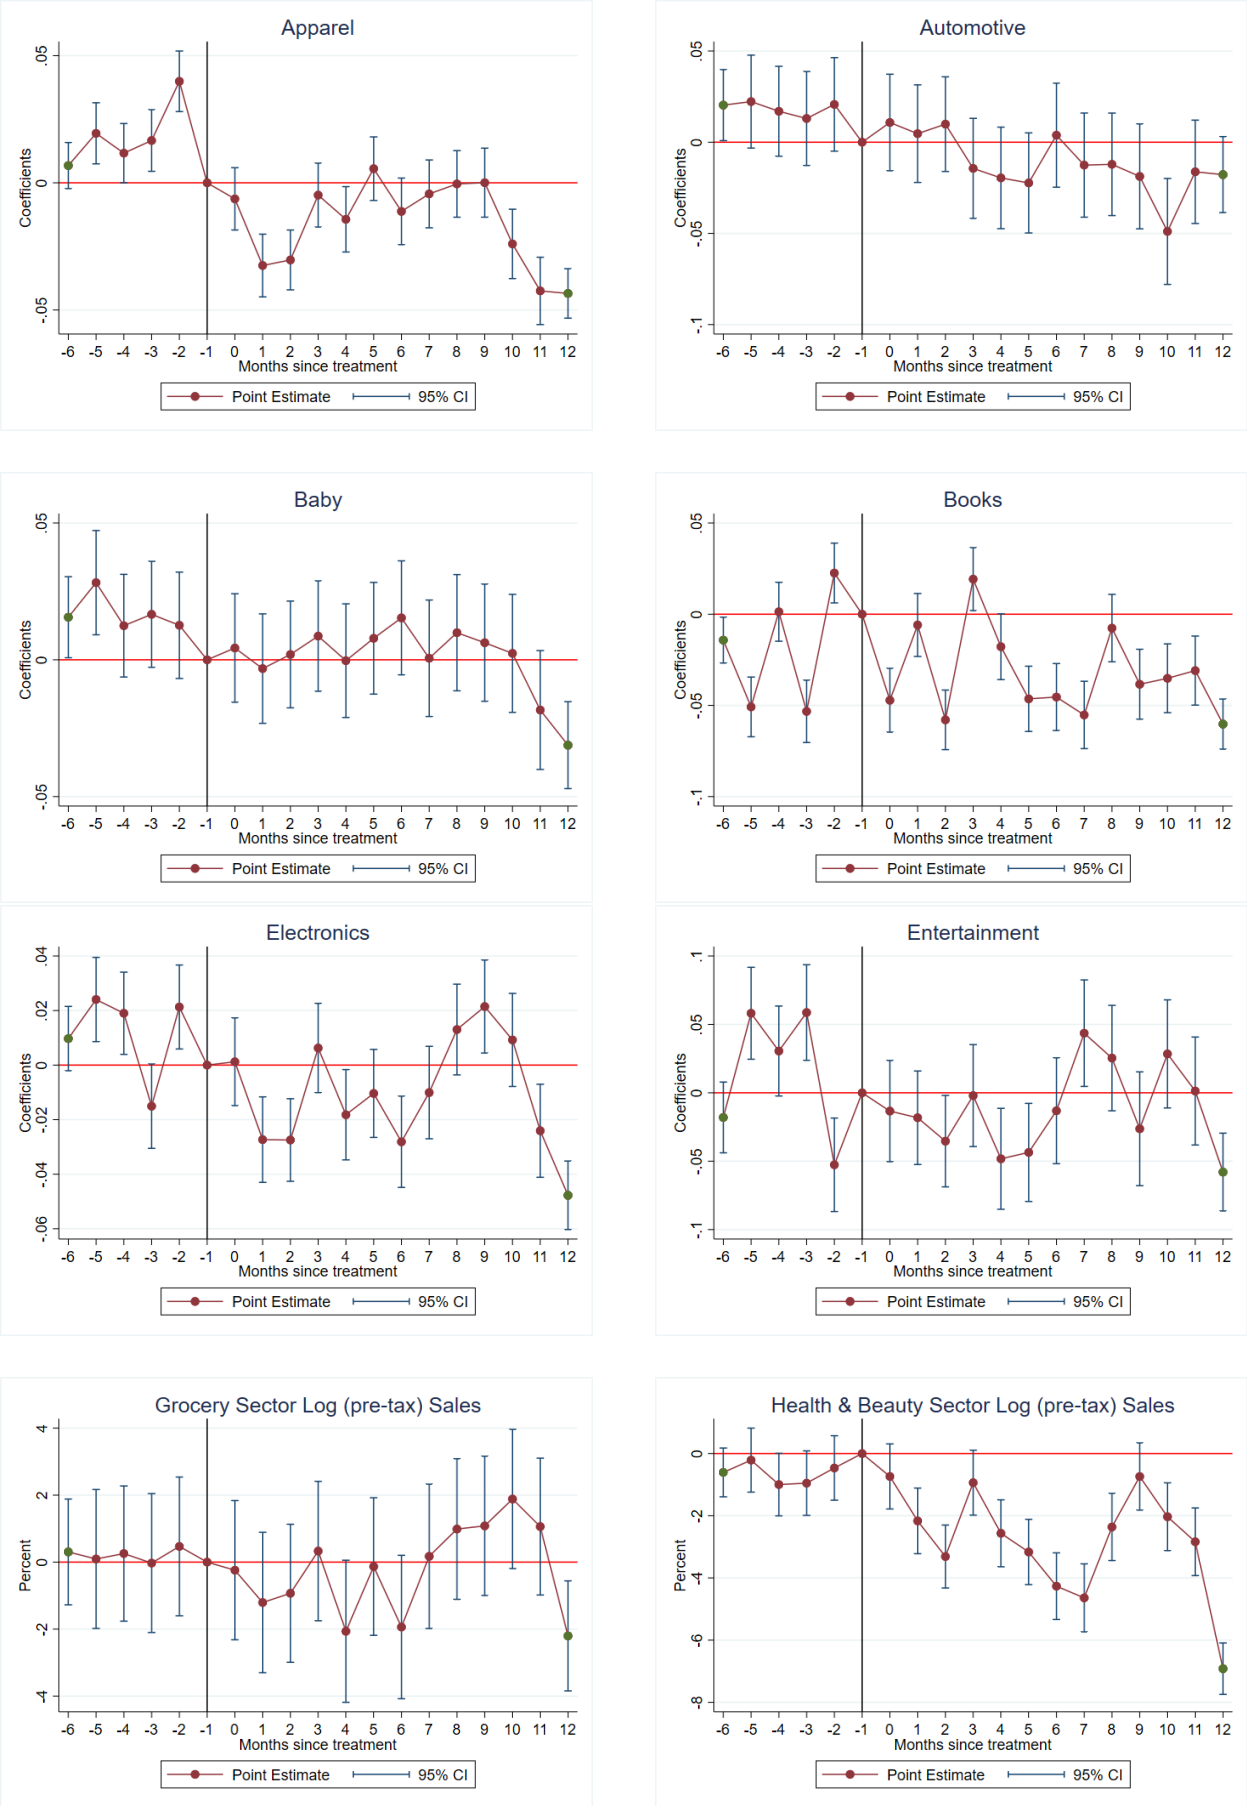

Figure A.13: (Cont.)Wayfair Effects over all sectors

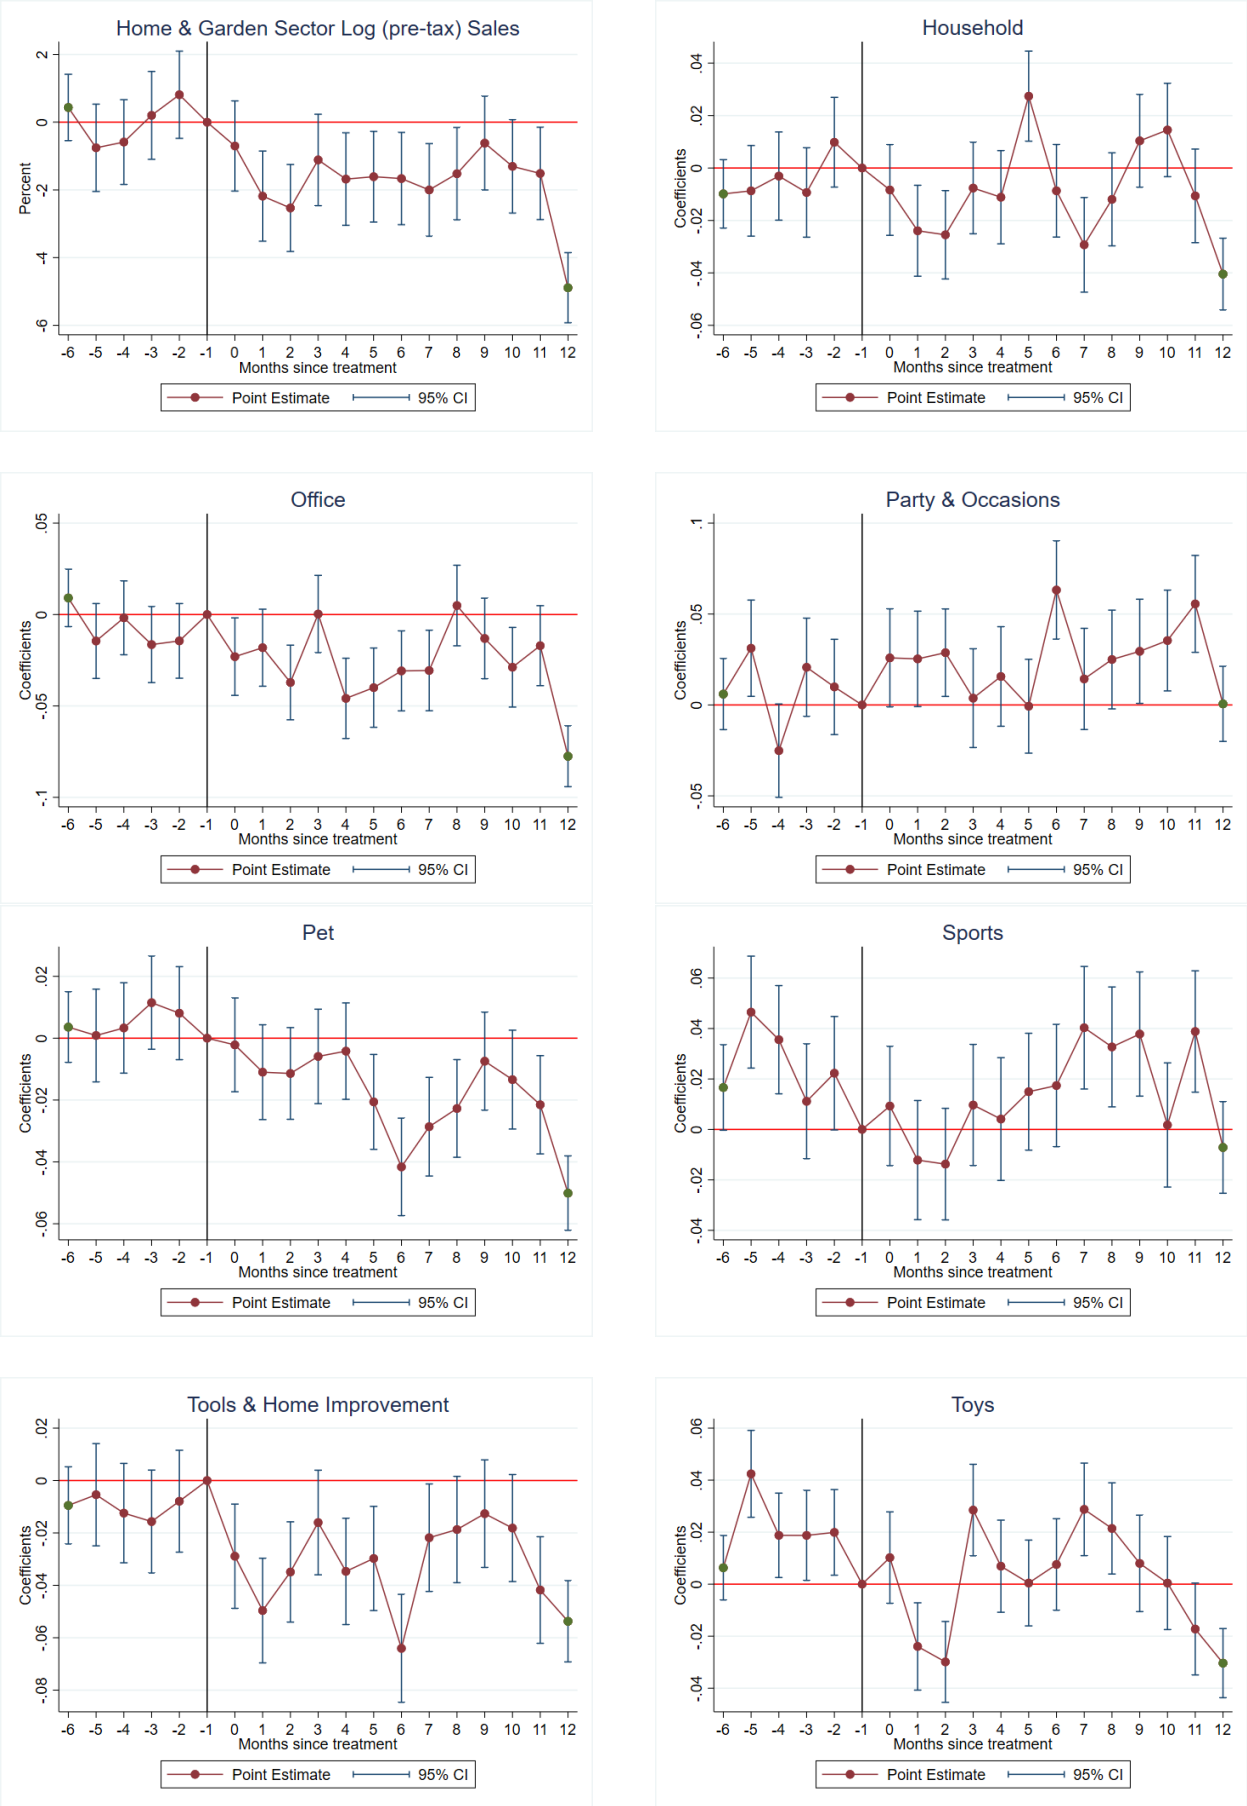

A.9 Assortment Responses

Figure A.14: Assortment Responses of the Untreated Retailers

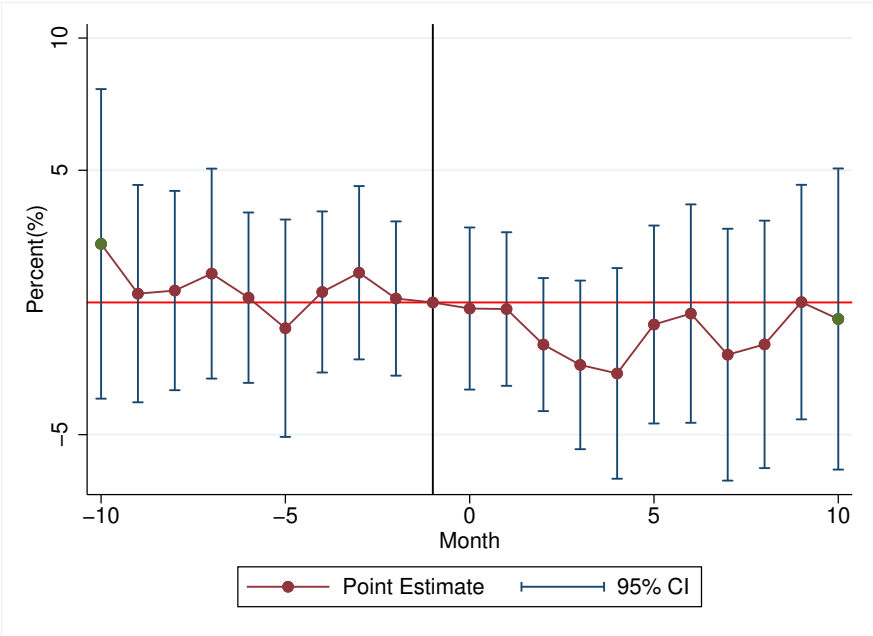

Note: Coefficients plots on β_k of our event study regression on the log of pre-tax sales as in (4). We run the regression on firm-channel-state level pre-tax prices and the number of unique items sold on leads and lags of the tax shock, controlling for firm-channel-time and firm-channel-state fixed effects.

A.10 Directed Acyclic Graphics for Causal Inference

We present the DAG for causal inference in the following graph. We are interested in the causal relationship between prices and mean utility, as is shown in the arrow from P_1 to Y_1. Here, unobserved demand shifter U_1^D is a confounding factor for causal inference. If firms conduct pricing to markets as Figure A.15a, we can use P_2 as an instrument for P_1, as long as P_2 is not affected by U_1^D. This is the "Hausman Instrument" argument. However, this strategy will fail if firms conduct uniform pricing, as Figure A.15b.

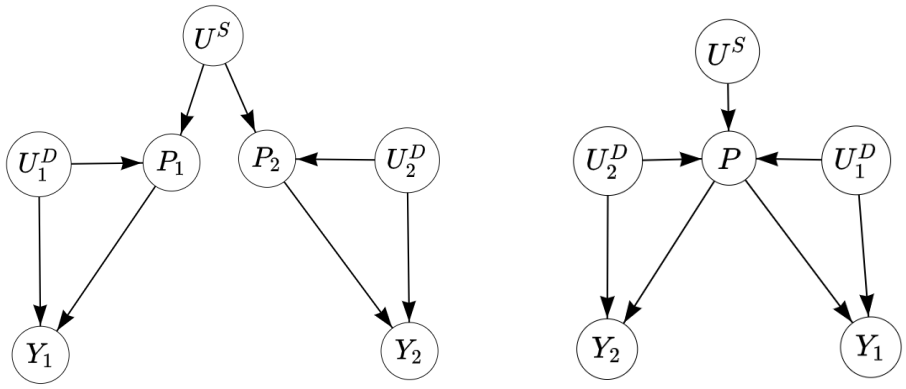

(a) Pricing to Market

(b) Uniform Pricing

If firms conduct uniform pricing, sales tax rate τ_1 can serve as an instrument for U_1^D, as in Figure A.16.

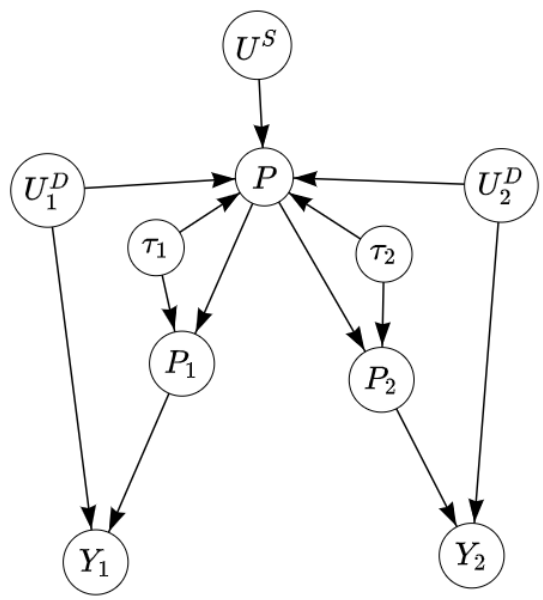

Figure A.16: tax shock level variations

A.11 Robustness Check Results

We present the robustness check results in this section.

A.11.1 Zero-Market Shares

Excluding zero-market-shares choices from consumers' choice sets may lead to underestimating the distance elasticities. Therefore, we estimate the within-chain distance elasticities using Poisson pseudo-maximum likelihood estimation (PPML). Table A.4 compares the PPML estimation with the implied number from our baseline estimation.

Table A.4: PPML Estimation Results

Group	IV-GMM ($\theta_f \beta$)	PPML (γ)
Rich	-0.46	-1.09 (0.007)
Poor	-0.51	-1.17 (0.006)

Note: We estimate the within-chain distance elasticities using Poisson pseudo-maximum likelihood estimation (PPML), and compare that with the implied distance elasticities in our baseline.

A.11.2 COVID shock and online-shopping

We test whether the distance elasticities are stable over time.

Table A.5: Distance Elasticities Over the Years

	Average	2017	2018	2019	2020	2021
β^t_{rich}	-1.09	-1.02	-1.08	-1.15	-1.12	-1.10
	(0.007)	(0.015)	(0.015)	(0.016)	(0.018)	(0.017)
β^t_{poor}	-1.17	-1.06	-1.12	-1.22	-1.21	-1.23
	(0.06)	(0.013)	(0.012)	(0.013)	(0.014)	(0.014)

Note: We estimate the within-chain distance elasticities for different years using PPML.

B. Theory Appendix

B.1 Micro-foundations of price index construction in descriptive analysis

We now discuss the measurement of $\Delta \ln P_{c,t}^{\text{on,T}}$ and the construction of the Wayfair Decision exposure as instruments. The change in price index can be approximated by the sum of price changes of each seller weighted by expenditure share. $d \ln P = \sum_j \omega_j \ln p_j$. To see this, suppose the within-group demand is CES with the exact price index $P = \left(\sum \xi_j p^{1-\eta}\right)^{1/(1-\eta)}$, where η is the within-group substitution elasticity and ξ_j is the taste shifters. Total differentiation yields:

$$\begin{aligned} d \ln P &= \frac{1}{1-\eta} d \ln \left(\sum_j \xi_j p_j^{1-\eta} \right) \\ &= \sum_j \frac{\xi_j p_j^{1-\eta}}{\sum_j \xi_j p_j^{1-\eta}} \frac{d p_j}{p_j} \\ &= \sum_j \omega_j d \ln p_j \end{aligned}$$

In our setting, the final price is the after-tax price $p_j = p_j^0 (1 + \tau_j)$. Therefore, the change in the price index of the treated group becomes:

$$d \ln P_{c,t}^{\text{on,T}} = \underbrace{\sum_j \omega_{j,c,t} d \ln p_{j,c,t}^0}_{\text{price changes}} + \underbrace{\sum_j \omega_{j,c,t} d \tau_{j,c,t}}_{\text{Wayfair exposure}}$$

B.2 Proof of Proposition 1

Our results closely follow McFadden (1980) and Feenstra (1995). It will be useful to state it here.

Theorem 1 *(Mcfadden 1980, Feenstra 1995) Let H be a non-negative function that satisfies a) Homogeneous of degree one. b) $H \to \infty$ as any of its arguments approaches infinity. c) The mixed partial derivatives of H w.r.t k variables exist and are continuous, non-negative if k is odd and non-positive if k is even, k = 1,..., N. Define the generalized extreme value distribution as:*

$$F(\epsilon_1, ... \epsilon_N) \equiv \exp(-H(e^{-\epsilon_1}, ..., e^{-\epsilon_N}))$$

Then the expected value of consumer utility (up to a constant) is given by the aggregate utility function

$$G(\delta_1, ..., \delta_N) \equiv \ln H(e^{\delta_1}, ... e^{\delta_N}) \tag{B.1}$$

and the choice probabilities can be obtained as

$$P_j = \frac{\partial G}{\partial \delta_j}$$

For our specifications,

$$H(\boldsymbol{x}) = \left(\sum_{j \in \text{on}} \left(\lambda \mathbb{1}_{C(j)=\text{on}} x_j \right)^{\theta_{on}} \right)^{\frac{1}{\theta_{on}}} + \sum_{f \in F} \left(\sum_{f} \left(\left(1 - \lambda \mathbb{1}_{C(j)=\text{on}}\right) x_j \right)^{\theta_f} \right)^{\frac{1}{\theta_f}}$$

First, since $H(k\boldsymbol{x}) = kH(\boldsymbol{x})$, it is homogeneous of degree one. We can also check b) and c) in Theorem 1 satisfies. Therefore, the proposed H function is suitable to generate GEV distribution. We now derive the choice probability P_j.

$$\begin{aligned}
P_j = \frac{\partial G}{\partial \delta_j} &= \frac{1}{H} \frac{\partial H}{\partial x_j} \frac{\partial x_j}{\partial \delta_j} \\
&= \frac{1}{H} \left(\lambda \mathbb{1}_{C(j)=\text{on}} x_j \right)^{\theta_{on}} \left(\sum_{j \in \text{on}} \left(\lambda \mathbb{1}_{C(j)=\text{on}} x_j \right)^{\theta_{on}} \right)^{\frac{1}{\theta_{on}} - 1} \\
&+ \frac{1}{H} \left(\left(1 - \lambda \mathbb{1}_{C(j)=\text{on}}\right) x_j \right)^{\theta_f} \left(\sum_{f} \left(\left(1 - \lambda \mathbb{1}_{C(j)=\text{on}}\right) x_j \right)^{\theta_f} \right)^{\frac{1}{\theta_f} - 1}
\end{aligned}$$

Then substitutes $x \equiv e^{\delta}$ gives us the market share function in (12). Since we already specified indirect utility, according to Roy's identity, the expected quantity is:

$$c_j = -\frac{\partial V_j / \partial \tilde{p}_j}{\partial V_j / \partial y} = -y \frac{\partial V_j}{\partial \delta_j} \frac{\partial \delta_j}{\partial \tilde{p}_j} = -\frac{\alpha y}{\tilde{p}_j} P_j \tag{B.2}$$

Note that here $\tilde{p}_j$ is the after-tax prices consumers face. From (B.2) we can solve for total (after-tax) sales revenue $S_j = c_j \tilde{p}_j$. Therefore, the sales revenue market share can be obtained

$$s_j = \frac{c_j \tilde{p}_j}{\sum_{j'} c_{j'} \tilde{p}_{j'}} = P_j$$

The expected maximized utility directly follows (B.1).